CodeSport

Sandya Mannarswamy

Welcome to CodeSport! In this issue, we will discuss the solutions to some of the challenging questions featured in last month's special New Year edition.

Many readers have sent in their answers to the questions we featured in the last edition of *LFY*. Some have requested that we keep the challenging questions open for the readers to come up with the solutions themselves. So let's look at a few of the questions from last month's column and discuss approaches to solving them instead of providing the complete solutions themselves. Congratulations to our readers Madhav Kumar, Vinay Saxena and Ravi Kiran for getting the correct answers to the maximum number of questions.

Minimum and maximum weighted cycles

To recall our first question from last month's column: given an undirected graph G (V, E) where V represents the set of vertices and E represents the set of edges, and a weight function c(u, v) that associates a non-negative weight with the edges of the graph, we were asked to come up with an algorithm to find the minimum weighted cycle.

We had discussed Djikstra's algorithm to find the shortest path (minimum weighted) between two vertices in a given graph. Can we somehow use Djikstra's algorithm to find the minimum weighted cycle involving an edge (u,v) in the graph? Recall that in order for the edge (u, v) to be part of a cycle, there should be another path between the vertices (u, v) which does not involve the edge (u, v). We can use Djikstra's shortest path to find the shortest path between vertices (u, v). Let us denote this shortest path as P2, whereas the direct path from vertex 'u' to vertex 'v' using the edge (u, v) is denoted as P1. The minimum weight of the cycle which involves the edge (u, v) is the sum of values P1 and P2. For each edge e ∈ E, we can compute the minimum weight cycle C(e) involving that edge. And the minimum over all such cycles is the minimum weighted cycle of the graph. I leave it to the reader to come up with the complete solution.

The second part of the problem asked readers to come up with an algorithm to find the maximum weighted cycle in the graph. Recall that there is no polynomial time algorithm to find the longest (maximum weighted) path between two vertices in a given directed graph. Can you extend this argument to the problem of finding the maximum weighted cycle in a graph using a formulation similar to the minimum weighted cycle we discussed earlier?

Finding the smallest 20 numbers in an array

To recap the second problem from last month's column: given an array of N unsorted integers, you were asked to find the smallest 20 integers from that array. The hint given was to use the linear time algorithm for median finding. Note that the question asks you to find the 20 smallest integers from the array but you do not have to order these 20 numbers from the smallest to the largest. You can find the median in linear time. Now if you partition the array using the median (remember the partition procedure from the quick sort algorithm the we discussed in a previous column?), you know that all elements partitioned below the median are smaller than the median and all elements partitioned above the median are larger than the median. Assuming that N is 31, if you find the median and partition around the median, you know that N/2 elements are smaller than the median; therefore, 15 elements which are partitioned below the median are smaller than the median. You need to find the remaining four smallest integers from the upper side of the partition by repeating the same procedure of finding the median and partitioning using the median. I leave it to you to use this idea to arrive at the complete solution.

Finding the 'k' minimum elements

Another question posted in last month's column was the problem of doing 'k' extract minimum operations repeatedly on a given array of N unsorted integers. The question assumes that the value of k is not known in advance. The naive solution is to sort the array in O (NlogN) and then extract the first k elements in the sorted array. However, the first part of the problem was to come up with an algorithm that can do this in O (Nk) time complexity. You know that finding the minimum in a given array of N integers can be done in O (N) and you can repeatedly perform 'find_minimum' k times to arrive at a solution whose complexity is O (Nk). The second part of the question was slightly tougher. It asked readers to find an algorithm for this problem whose complexity is O (N) + klogN? How do you do this?

The trick is to recognise that you perform 'k' extract minimum operations. Note that in order for the overall complexity to be O (N) + klogN, each of the extract minimum should be taking O (logN). What is the data structure that performs extract minimum with a complexity of O (logN)? Of course, it is the binary heap. You also know that you can build a binary heap in O (N) time. Therefore, the solution is to build a binary heap from the array of N unsorted integers and perform 'k' *extract_min* operations on the binary heap.

RAIL: Resource Acquisition Is Initialisation

This is a popular design pattern used in C++ in order to ensure the release of resources acquired at the beginning of a function scope. Resources acquired in a particular scope should be released before leaving the scope, unless the ownership is being transferred to another scope or object explicitly so that the new owner will take the responsibility of releasing the resource after using it. Typically, this is achieved by having a resource acquire call at the beginning of the function (where the resource is going to be used) and a resource release call at the end of the function (after which you do not need the resource anymore). The resources can be memory, file, lock, etc. Very often you see a pair of function calls such as (*malloc, free*) (*mutex_lock, mutex_unlock*) (*fopen, fclose*) at the beginning and end of the function.

For a programmer it is sometimes possible to forget writing the code that performs the resource release. If there are many paths of control from which the function can return including some paths where you return due to errors, the programmer can forget releasing the resource along one of these paths. This can result in resource leakage.

Consider the following piece of code:

```
bool foo ()
{
  char * mystr   new char [100];
  if (cond1)
    do something //with acquired memory
  else if (cond2)
      return false
  else
    throw "ERROR"

  // This may not be invoked in all return paths
  // hence memory can be leaked
  delete [] mystr
  return true;
}
```

In the above snippet of code, there are control flow paths in the routine where the programmer has forgotten to free the memory acquired before returning from the function. This is a common programming error often leading to memory leaks.

Consider the following piece of code:

```
bool bar ()
{
      pthread_mutex_lock(&mutex1)
      if (cond1)
          [illegible]
      [illegible])
          return false;
    else
          throw "ERROR"

    // The lock release may not be invoked in all return path s
    // hence the code end up in a deadlock
    pthread_mutex_unlock(&mutex1)
    return true;
}
```

In the above code snippet, there are paths of control flow where the programmer has forgotten to release the lock, which can lead to a possible deadlock. This is again a common programming error when the flow of control in the program is quite complex, as in the example above.

RAIL is a programming idiom/design pattern aimed at solving the problem of a missing resource release along all return paths from a given scope. The idea is to wrap the resource acquire in the constructor of an object whose scope is the same as the function scope. The destructor of the object contains the call to the resource release function. Since the destructor of the object will be called whenever it goes out of scope automatically, the resource release function will be called automatically whenever there is a control flow path returning from the function.

Note that it is not mandatory that the resources need to be acquired in the constructor of the object, though it is often convenient to do so. However the key to RAIL is that the resources get released in the destructor of the object. Hence, RAIL is also sometimes known as *Resource Release is Finalisation* idiom. I leave it to the readers to come up with code that uses RAIL to avoid the resource leakage in the two examples of code snippets shown earlier. The resource release functions we have seen are 'free' and 'mutex_unlock' which do not throw exceptions. What would happen if the resource release function ends up throwing an exception? How would RAIL handle such cases?

I would like to request readers to send in their solutions to the remaining questions from last month's column. Also, if you are preparing for computer science programming interviews, you may find it useful to visit my programming interviews discussion group 'Computer Science Interview Training (India)' on LinkedIn.

If you have any favourite programming puzzles that you would like to discuss on this forum, please send them to me along with your solutions and feedback at sandyasm_AT_yahoo_DOT_com. Till we meet again next month, happy programming! END

About the author:

Sandya Mannarswamy The author is a specialist in compiler optimisation and works at Hewlett Packard India. She has a number of publications and patents to her credit, and her areas of interest include virtualisation technologies and software development tools.

A Voyage to the Kernel

Part 21

Segment: 4.1, Day 20

In Segment 3, we'd looked at some fundamental aspects of Linux kernel module programming and we discussed commands like *lsmod* and files like */etc/modprobe.conf*. We also meddled with *__init* and *__exit* while writing a simple module. This new segment is dedicated to kernel programming, in general.

The easiest way to add code to the kernel is to just add the lines to the source tree and recompile it. You may recall that you can actually pre-configure the kernel you are about to compile (please refer to the previous segment for more details).

We have also seen that we can add some code to the kernel even when it is alive. This is what we did when we wrote simple kernel modules (or loadable kernel modules) and loaded them to the kernel. Typically, if you just want to write a driver for new hardware that you have made, and enable the device, then you don't have to compile the entire kernel to make the device run. A module will serve the purpose. So is the case with system calls. This is one of the properties of the Linux kernel. Well, for our discussions, I'm going to call the kernel that gets loaded while booting up as the 'base kernel' and loadable ones as modules (as a generic term).

Module programming became popular during the mid-90s and attracted many developers (and vendors). But the earlier modules were not that easily 'loadable'. It was only by the beginning of this decade that they attained the shape of modules, as we see them today.

Why should it be loadable?

This is one of the fundamental questions that beginners ask. "Why should I write a loadable module? Why not incorporate it into the kernel code itself?" Well, you may find some reasons if you think of yourself as a hardware vendor who's going to offer a new piece of hardware. As discussed before, the user's convenience is a factor. They don't have to rebuild the kernel each time they add a new device driver.

Also, if you try to rebuild the kernel with your code, then it may be hard for the developers to fix the error that may pop up during the process. Sometimes the error may even prevent you from logging into the system if you have a serious bug in the code you just added (since it will be executed during the boot stage).

If you look at the resource side, then you can see that the module saves much of your memory as you will be loading it only when it is needed. Another fact is that developers (say, hardware vendors) find it easy to maintain their code if it is written as a module.

This is even helpful during the development phase. You can experiment with your code very easily if you write it as a module. It will save you time as you don't have to reboot the system each time you update the code. But you also need to note that if you are writing a driver for a device that is essential for providing the service, then you should try to incorporate the code into the base kernel itself.

Useful utilities

While developing the module, you may find the following utilities quite handy:

- *insmod* – to load a module into the kernel (refer to the previous segment—for your information, I'm putting the *moduleloader* header file code here):

```
#ifndef _LINUX_MODULELOADER_H
#define _LINUX_MODULELOADER_H
#include <linux/module.h>
#include <linux/elf.h>
int module_frob_arch_sections(Elf_Ehdr *hdr,
                              Elf_Shdr *sechdrs,
                              char *secstrings,
                              struct module *mod);
unsigned int arch_mod_section_prepend(struct module *mod,
unsigned int section);
void *module_alloc(unsigned long size);
void module_free(struct module *mod, void *module_region);
int apply_relocate(Elf_Shdr *sechdrs,
           const char *strtab,
           unsigned int symindex,
                     struct mod
int apply_relocate_add(Elf_Shdr *sechdrs,
                       const char *strtab,
                       unsigned int symindex,
                       unsigned int relsec,
                       struct module *mod);
int module_finalize(const Elf_Ehdr *hdr,
                    const Elf_Shdr *sechdrs,
                    struct module *mod);
void module_arch_cleanup(struct module *mod);
#endif
```

- *rmmod* – remove a module from the kernel
- *depmod* – to find out the dependencies between modules
- *kerneld* – for invoking the *kerneld* daemon
- *ksyms* – To display the symbols exported by the kernel which the module can use
- *lsmod* – To list the loaded modules (see the code for the structure of its output)

```
aasaevinayak@GNU BOX:~$ lsmod
Module                  Size  Used by
rfcc                    3516  1052
aes_i586                8124  1053
aes_generic            27484  1 aes_i586
ecb                     2524  1
binfmt_misc             8356  1
ppdev                   6068  0
videobuf                78760  0
videodev                85268  0
v4l1_compat            121608  1 videodev
des_crypt               12828  0
<<<CODE TRUNCATED>>>
```

- *modinfo* – To display the contents from the *modinfo* section. For example, the information about my video device can be displayed using the following code:

```
aasaevinayak@GNU BOX:~$ modinfo videodev
filename:       /lib/modules/2.6.31-17-generic-pae/kernel/drivers/media/
video/videodev.ko
alias:          char-major-81-*
license:        GPL
description:    Device registrar for Video4Linux drivers v2
author:         Alan Cox, Mauro Carvalho Chehab <mchehab@infradead.org>
srcversion:     943A577655280E3A6C88777
depends:        v4l1-compat
vermagic:       2.6.31-17-generic-pae SMP mod_unload modversions 586
```

- *modprobe* – helps you to load/remove modules in a better way (We will discuss this later in this segment as well.) Here is the output that shows how *dummy.ko* is loaded into the kernel.

```
aasaevinayak@GNU BOX:~$ strace modprobe dummy
execve("/sbin/modprobe", ["modprobe", "dummy"], [/* 35 vars */]) = 0
brk(0)                                  = 0x8dee000
access("/etc/ld.so.nohwcap", F_OK)      = -1 ENOENT (No such file or
directory)
mmap2(NULL, 8192, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_
ANONYMOUS, -1, 0) = 0xb770a000
access("/etc/ld.so.preload", R_OK)      = -1 ENOENT (No such file or
directory)
open("/etc/ld.so.cache", O_RDONLY)      = 3
fstat64(3, {st_mode=S_IFREG|0644, st_size=68604, ...}) = 0
mmap2(NULL, 68604, PROT_READ, MAP_PRIVATE, 3, 0) = 0xb7760000
close(3)                                = 0
access("/etc/ld.so.nohwcap", F_OK)      = -1 ENOENT (No such file or
directory)
open("/lib/tls/i686/cmov/libc.so.6", O_RDONLY) = 3
read(3, "\177ELF\1\1\1\0\0\0\0\0\0\0\0\0\3\0\3\0\1\0\0\0\300\1\0\
0004\0\0\0"..., 512) = 512
fstat64(3, {st_mode=S_IFREG|0755, st_size=1319364, ...}) = 0
mmap2(NULL, 1325416, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_
DENYWRITE, 3, 0) = 0xb7618000
mmap2(0xb7760000, 12288, PROT_READ|PROT_WRITE, MAP_
PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x15e) = 0xb7760000
mmap2(0xb7763000, 10680, PROT_READ|PROT_WRITE, MAP_
PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0xb7763000
close(3)                                = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_
ANONYMOUS, -1, 0) = 0xb7761000
set_thread_area({entry_number:-1 -> 6, base_addr:0xb75b16c0,
limit:1048575, seg_32bit:1, contents:0, read_exec_only:0, limit_in_pages:1,
seg_not_present:0, useable:1}) = 0
mprotect(0xb7760000, 8192, PROT_READ)   = 0
mprotect(0x8062000, 4096, PROT_READ)    = 0
mprotect(0xb772a000, 4096, PROT_READ)   = 0
munmap(0xb7760000, 68604)               = 0
uname({sys="Linux", node="GNU BOX", ...}) = 0
fstat64(2, {st_mode=S_IFCHR|0620, st_rdev=makedev(136, 0), ...}) = 0
brk(0)                                  = 0x8dee000
brk(0x8e0f000)                          = 0x8e0f000
open("/etc/modprobe.conf", O_RDONLY)    = -1 ENOENT (No such file or
directory)
open("/etc/modprobe.d", O_RDONLY|O_NONBLOCK|O_LARGEFILE|O_
DIRECTORY|O_CLOEXEC) = 3
fcntl64(3, F_GETFD)                     = 0x1 (flags FD_CLOEXEC)
getdents(3, /* 12 entries */, 32768)    = 380
getdents(3, /* 0 entries */, 32768)     = 0
close(3)                                = 0
open("/etc/modprobe.d/alsa-base.conf", O_RDONLY) = 3
fstat64(3, {st_mode=S_IFREG|0644, st_size=2497, ...}) = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_
ANONYMOUS, -1, 0) = 0xb772b000
read(3, "# autoloader aliases\ninstall sou"..., 4096) = 2497
read(3, "", 4096)                       = 0
close(3)                                = 0
munmap(0xb772b000, 4096)                = 0
```

I also suggest that beginners to take a look at the */proc/module* (please note that this will get updated even if you leave it open):

```
marlee 2836 0 - Live 0xf858d000
hid 51492 1 usbhid, Live 0xf8c5e000
olc 3516 1163 - Live 0xfa37d000
```

```
aes_i586 8124 1 - Live 0xf86a6000
aes_generic 27484 1 aes_i586, Live 0xfa360000
ecb 2524 1 - Live 0xf8501000
binfmt_misc 8868 1 - Live 0xf8566000
ppdev 6068 0 - Live 0xf8578000
vboxnetadp 78760 0 - Live 0xfa06a000
vboxnetflt 86288 0 - Live 0xfa03a000
vboxdrv 121908 1 vboxnetflt, Live 0xf9ff8000
dm_crypt 12938 0 - Live 0xf8822000
bridge 47952 0 - Live 0xf9753000
stp 2272 1 bridge, Live 0xf971c000
bnep 13080 2 - Live 0xf9713000
snd_hda_codec_idt 58844 1 - Live 0xf96c4000
snd_hda_intel 26884 4 - Live 0xf9690000
snd_hda_codec 76708 3 snd_hda_codec_idt,snd_hda_intel, Live 0xf8874000
snd_hwdep 7200 1 snd_hda_codec, Live 0xf8863000
```

Another important file is */etc/modprobe.d/blacklist*. You may also check other files in this folder. This file (which contains a list) allows you to block some of the modules from getting loaded, thus paving the way to use alternative drivers. Sometimes a fake module may also try to block your 'free access' to a device. You can use this file to block the same from getting loaded. Here is a sample of an entry for your reference.

```
# evbug is a debug tool that should be loaded explicitly
blacklist evbug
# these drivers are very simple, the HID drivers are usually preferred
blacklist usbmouse
blacklist usbkbd

# replaced by e100
blacklist eepro100

# replaced by tulip
blacklist de4x5
# causes no end of confusion by creating unexpected network interfaces
blacklist eth1394
# snd_intel8x0m can interfere with snd_intel8x0, doesn't seem to support much
# hardware on its own (Ubuntu bug #2011, #6810)
blacklist snd_intel8x0m

<<<OUTPUT TRUNCATED>>>
```

Having discussed these aspects, let's write a module that does more functions than the first one we wrote (refer to Segment 3 for the first tutorial):

Open your editor and write the following lines of code:

```
#include <linux/module.h>
#include <linux/moduleparam.h>
#include <linux/kernel.h>
#include <linux/init.h>
#include <linux/stat.h>

MODULE_LICENSE("GPL v2");
MODULE_AUTHOR("Aasis Vinayak PG");

static short int short_integer = 0;
static int integer = 100;
static long int long_integer = 999999;
static char *string_entry = "GNU Linux";
static int integerArray[2] = { 1, 1 };
static int arr_arg = 0;

module_param(short_integer, short, S_IRUSR | S_IWUSR | S_IRGRP |
S_IWGRP);
MODULE_PARM_DESC(short_integer, "A short integer");
module_param(integer, int, S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH);
MODULE_PARM_DESC(integer, "An integer");
module_param(long_integer, long, S_IRUSR);
MODULE_PARM_DESC(long_integer, "A long integer");
module_param(string_entry, charp, 0000);
MODULE_PARM_DESC(string_entry, "A string");
module_param_array(integerArray, int, &arr_arg, 0000);
MODULE_PARM_DESC(integerArray, "An array of integers");

static int __init tutorial1_init(void)
{
	int i;
	printk(KERN_INFO "Welcome to another day of this Voyage\n");
	printk(KERN_INFO "short_integer is a short integer: %hd\n", short_integer);
	printk(KERN_INFO "integer is an integer: %d\n", integer);
	printk(KERN_INFO "long_integer is a long integer: %ld\n", long_integer);
	printk(KERN_INFO "string_entry is a string: %s\n", string_entry);
	for (i = 0; i < (sizeof integerArray / sizeof (int)); i++)
	{
		printk(KERN_INFO "The integerArray[%d] = %d\n", i, integerArray[i]);
	}
	printk(KERN_INFO "has %d arguments.\n", arr_arg);
	return 0;
}

static void __exit tutorial1_exit(void)
{
	printk(KERN_INFO "Exiting the module\n");
}

module_init(tutorial1_init);
module_exit(tutorial1_exit);
```

As you can see in the code, after including the relevant header files (necessary for compiling) and declaring the variables, we start the kernel specific code portion. The next part of the code is used for handing the parameter. And we have already seen why we use *_init* and *_exit*. Even if you are a beginner, I am sure that you would have guessed what this module is doing—we have written a module that can receive values as parameters!

Now we need to write a *make* file. And here is the code for it.

```
[illegible]
[illegible]
	make -C /lib/modules/$(shell uname -r)/build M=$(PWD) modules
clean:
	make -C /lib/modules/$(shell uname -r)/build M=$(PWD) clean
```

Save this in the same directory in which you have saved your C program. Now you can *cd* to that directory from the terminal and issue the *make* command:

```
aasisvinayak@GNU-BOX:~/Desktop/modules$ make
make -C /lib/modules/2.6.31-17-generic-pae/build M=/home/aasisvinayak/Desktop/modules modules
make[1]: Entering directory `/usr/src/linux-headers-2.6.31-17-generic-pae'
  CC [M]  /home/aasisvinayak/Desktop/modules/VTK_module_tutorial_1.o
  Building modules, stage 2.
  MODPOST 1 modules
  CC      /home/aasisvinayak/Desktop/modules/VTK_module_tutorial_1.mod.o
  LD [M]  /home/aasisvinayak/Desktop/modules/VTK_module_tutorial_1.ko
make[1]: Leaving directory `/usr/src/linux-headers-2.6.31-17-generic-pae'
```

You can now see a few other files created in the directory (refer to Figure 1).

If you are a Java developer, then you must have used *getters* and *setters* while meddling with Java beans. Here, in one of the automatically generated files (source), you can find something similar:

```
#include <linux/module.h>
#include <linux/vermagic.h>
#include <linux/compiler.h>

MODULE_INFO(vermagic, VERMAGIC_STRING);

struct module __this_module
__attribute__((section(".gnu.linkonce.this_module"))) = {
 .name = KBUILD_MODNAME,
 .init = init_module,
#ifdef CONFIG_MODULE_UNLOAD
 .exit = cleanup_module,
#endif
 .arch = MODULE_ARCH_INIT,
};

static const struct modversion_info ____versions[]
__used
__attribute__((section("__versions"))) = {
	{ 0x4eaaf241, "module_layout" },
	{ 0xb2242be2, "param_get_short" },
	{ 0x4f555eadb, "param_set_short" },
	{ 0x89980fe81, "param_get_int" },
	{ 0xff964b26, "param_set_int" },
	{ 0x8ec8b063, "param_get_long" },
	{ 0x5657cb88, "param_set_long" },
	{ 0x41344088, "param_get_charp" },
	{ 0x6ad96514, "param_set_charp" },
	{ 0x45eab6c3, "param_array_get" },
```

Figure 1: More files are created after running make

```
	{ 0x40847727, "param_array_set" },
	{ 0xb72397d5, "printk" },
};

static const char __module_depends[]
__used
__attribute__((section(".modinfo"))) =
"depends=";

MODULE_INFO(srcversion, "FEC7A3AF80AAF8E883EA05E");
```

Though we have seen a few aspects/properties of these files already we will revisit these areas and cover them in depth. Since this is an introductory column for the new segment, I'm avoiding the minute details.

Now you can load the module to the kernel and check if it is getting loaded properly:

```
aasisvinayak@GNU-BOX:~/Desktop/modules$ sudo insmod VTK_module_tutorial_1.ko
aasisvinayak@GNU-BOX:~/Desktop/modules$ lsmod
Module                  Size  Used by
VTK_module_tutorial_1   2816  0
rndis                   7836  0
fat                    51492  1 msdos
<<<OUTPUT TRUNCATED>>>
```

Now let's test it by passing a parameter. We will pass a valid argument at first and then we will pass an invalid one.

```
aasisvinayak@GNU-BOX:~/Desktop/modules$ sudo insmod VTK_module_tutorial_1.ko string_entry="hooray" integerArray=-1
aasisvinayak@GNU-BOX:~/Desktop/modules$ sudo rmmod VTK_module_tutorial_1aasisvinayak@GNU-BOX:~/Desktop/modules$ sudo insmod VTK_module_tutorial_1.ko string_entry="hoo" integerArray=-1 invalid_entry="null"
insmod: error inserting 'VTK_module_tutorial_1.ko': -1 Unknown symbol in module
```

We have not yet covered some other basic elements like checking the modules (loaded) regularly and removing (cleaning) the ones that are not in use. We will cover these topics in the next segment.

Happy kernel hacking!

By: Aasis Vinayak PG

The author is a hacker and a free software activist who does programming in the open source domain. He is the developer of V language—a programming language that employs AI and ANN. His research work/publications are available at www.aasisvinayak.com

Powered by
naukri.com
India s No 1 Job Site

Top Jobs

Post: **Linux Administrator**
Company: Goldstone Technologies Ltd.
Profile: Strong experience in Linux is desirable
Exp: 7-12
Location: Bengaluru
Email:
dvshaileshkumar@goldstonetech com

Post: **Unix Administrator**
Company: Pinceton Information Pvt Ltd
Profile. Looking for a strong Unix professional with solid experience on Solaris, Veritas product suite Veritas cluster server and Legato backup
Exp: 4-9
Location: Chennai
Email
venkat rao@princetoninformation com

Post: **Sr Linux Developer**
Company: PT Pure Testing Software Pvt. Ltd
Profile: Should have 2-4 years experience in development with good expertise on Unux.
Exp: 2-4
Location: Noida
Email: amit@puretesting com

Post: **Unix Developer**
Company: GCTECH
Profile: Should be BE / BTech with 1-3 years experience in projects involving ADM activities on Unix / Linux platform.
Exp: 2 6
Location: Bengaluru
Email: info gctechnologies@gmail.com

Post: **System Administrator (SUSE Linux)**
Company: Serve all Enterprise Solutions Ltd
Profile: Ideal candidate should have about 4 years of hands-on experience in Linux (SUSE SLES 9/10) &amp AIX system administration
Exp: 4-8 **Location**: Bengaluru
Email: surya kondeti@servealltech com

Post: **HP-UX Engineer**
Company: Computer Centre Ltd.
Profile: Must have Diploma / Degree qualification in Computer Science or similar with Unix / Linux certifications
Exp: 5-10 **Location**: Tanzania
Email: careers cc@omnetac.com

Post: **Technical Support - Linux Professional**
Company: Libsys Technologies Pvt Ltd.
Profile: Should provide technical support on LibSys software on Windows / Linux / Unix to sites in their region through email / phone / visits
Exp. 2-7 **Location**: Gurgaon
Email: hm@libsys co in

Post: **Linux Administrator**
Company: ASM Technologies Ltd
Profile. Candidate should have 3-5 years of Unix support experience in a large production environment.
Exp: 3 5 **Location**: Bengaluru
Email: harsha ms@asmltd com

Post: **Unix Specialist**
Company: IBHSPL
Profile: Knowledge of wide range of enterprise infrastructure technology, software and business objectives is desirable.
Exp: 9 12 **Location**: Pune
Email: myitjobs@talentbucks.com

Post: **Software Developer - Linux, Unix**
Company: Xecute HR
Profile: Must be BE / BTech with 2 years relevant experience in programming (C or C++)
Exp: 2 5 **Location**: Bengaluru
Email: vandana rana@xecutehr com

Post: **Linux Production Support**
Company: B R Ray Soft Tech Pvt. Ltd
Profile: Should have 2+ years experience in Unix + Support project experience and knowledge / experience in Middlewares
Exp: 2-6, **Location**: Bengaluru
Email: balaji@braysoft.com

Post: **Unix Administratore**
Company: Yahoo Software Development India Pvt Ltd
Profile: Should have complete understanding of a product life cycle and be able to relate it with the business Performance Tuning (UNIX, Networking, Web).
Exp: 8 13
Location: Bengaluru
Email: lathav@yahoo-inc.com

Post: **AIX Administrator - Level 1**
Company: DCM Data Systems
Profile: Should have at least t 5+ years of experience in Unix administration (AIX/ Solaris, Linux/ HP-UX)
Exp: 3-8
Location: Bengaluru
Email. kiran@dcmds com

Post: **Linux Systems Engineer**
Company: Strategies Ingenious Software Solutions Pvt Ltd
Profile: Must have hands on-experience at Linux
Exp: 3-4
Location: Saudi Arabia
Email: job@stratiss com

Post: **Linux Administrator - Voice profile**
Company: Apollo Health Street Ltd
Profile: Any Science Graduate preferably MCA / B TECH / BE with at least 1 year of technical support experience in handling International client (US/UK) and Intermediate knowledge on Unix / Linux may apply
Exp: 1-3
Location: Delhi
Email:
hemanthkumar_k@apollohealthstreet.com

Post: **Network Administrator**
Company: mBit Infotech Pvt. Ltd
Profile: Candidate should have strong experience in in Linux flavour (Linux / Ubuntu / CentOS / RHEL)
Exp: 3 6, **Location**: Delhi
Email: career@mbit co in

LINUX For You

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

HIGHLIGHTS

- A cost effective marketing tool
- A user friendly format for customers to contact you
- A dedicated section with yellow back ground, and hence will stand out
- Reaches to tech-savvy IT implementers and software developers
- 80% of LFY readers are either decision influencers or decision takers
- Discounts for listing under multiple categories
- Discounts for booking multiple issues

FEATURES

- Listing is categorised on the basis of products and services
- Complete contact details plus 30 word description of organisation
- Option to print the LOGO of the organisation too (extra cost)
- Option to change the organisation description for listings under different categories

TARIFF

Category Listing	
ONE Category	Rs 2,000
TWO Categories	Rs 3,500
THREE Categories	Rs 4,750
ADDITIONAL Category	Rs 1,000

Value-add Options	
LOGO-plus Entry	Rs 500
Highlight Entry (white background)	Rs 1,000
Per EXTRA word (beyond 30 words)	Rs 50

KEY POINTS

- Above rates are per category basis.
- Above rates are charges for publishing in a single issue of LFY
- Max. No. of Words for Organisation Description: 30 words.

TERMS & CONDITIONS

- Fill the form (below)
- You can use multiple copies of the form for multiple listings under different categories
- Payment to be received along with booking

✂ Tear & Send Tear & Send ✂

ORDER FORM

Organisation Name (70 characters) ____________________

Description (30 words) ____________________

Email ____________________ Website: ____________________

STD Code: __________ Phone: ____________________ Mobile: ____________________

Address (will not be published) ____________________

____________________ City/Town: ____________________ Pin-code: __________

CATEGORIES

☐ CONSULTANTS
☐ CONSULTANT (FIRM)
☐ EMBEDDED SOLUTIONS
☐ ENTERPRISE COMMUNICATION SOLUTIONS
☐ HIGH PERFORMANCE COMPUTING
☐ IT INFRASTRUCTURE SOLUTIONS
☐ LINUX BASED WEB HOSTING
☐ MOBILE SOLUTIONS
☐ SOFTWARE DEVELOPMENT
☐ TRAINING FOR PROFESSIONALS
☐ TRAINING FOR CORPORATE
☐ THIN CLIENT SOLUTIONS

Please find enclosed a sum of Rs. __________ by DD/ MO//crossed cheque* bearing the No. ____________________ dt. __________ in favour of **EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs. 50 on non-metro cheque) towards the cost of __________ FOSS Yellow Pages advertisement(s)**

or charge my credit card ☐ VISA ☐ Master Card Please charge Rs. __________

against my credit card No. ☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐ C V V No. __________ (Mandatory)

Date of Birth ____ / ____ / ______ (dd/mm/yy) Card Expiry Date ____ / ____ (mm/yy)

Signature (as on the card)

EFY Enterprises Pvt Ltd D-87/1 Okhla Industrial Area, Phase 1, New Delhi 110 020
Ph: 011 26810601-03 Fax: 011 26817565 Email: info@efyindia.com Website: www.efyindia.com

To Book Your Listing, Call: Dhiraj (Delhi: 09811206582), Somaiah (B'lore: 09986075717)

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact
Dhiraj (Delhi) **09611206662,** Somaish (Bengalore) **09966075717**

Consultant (Firm)

IB Services

Free Installation of GNU/Linux on Laptops and Desktops. Thin client solutions based on Debian and Ubuntu. Laptops and Desktops pre-installed with Debian and Ubuntu. Migration to GNU/Linux. Data Recovery Navi Mumbai

Kerala
Mobile: 09847446080
Email: ibmanoj@gmail.com
Web: www.ibservices.in

OS3 Infotech

•Silver Solutions Partner for Novell •High Availability Computing Solutions •End-to-end Open Source Solutions Provider •Certified Red Hat Training Partner •Corporate and Institutional Training

Navi Mumbai
Mobile: 09324113579
Email: info@os3infotech.com
Web: www.os3infotech.com

Taashee Linux Services

100% Support on LINUX, OSS & JBOSS related projects. We specialize in high availability and high-performance clusters,remote and onsite system management maintenance services systems planning, Linux & JBOSS consulting & Support services.

Hyderabad
Mobile: 09392493753, Fax: 040-40131726
Email: niranjan.sunnar@taashee.com
Web: www.taashee.com

Torrid Networks Pvt Ltd

Torrid is a leading provider of Information technology consulting with focus on Information Security Services and Open Source Solutions.

Noida
Tel: 0120-4545100, Mobile: 09815685593
Fax: 0120-4235064
Email: tarun.kansare@torridnetworks.com
Web: www.torridnetworks.com

Computer (UMPC) For Linux And Windows

Comptek International

World's smallest computer comptek wibrain B1 UMPC with Linux,Touch Screen, 1 GB RAM 60GB, Wi-Fi, Webcam, upto 6 hour battery (opt.), USB Port, max 1600×1200 resolution, screen 4.8", 7.5 x3.25" Size, weight 526 gm.

New Delhi
Mobile: 09968756177 Fax: 011-26187551
Email: comptekdelhi@compteki.com
Web: www.compteki.com
or www.compteki.in

Education & Training

Aptech Limited

IT, Multimedia and Animation Education and Training

Mumbai
Tel: 022-28272300, 66462300
Fax: 022-28272399
Email: customercare@aptech.ac.in
Web: www.aptech-education.com, www.arena-multimedia.com

IT-Campus: Academy of Information Technology

IT training and solution company with over 12 years of experience. - RHCE •Software Training •Hardware Training •Multimedia And Animation •Web Designing •Financial Accounting

Kota (Raj.)
Tel: 0744-2503155, Mobile: 09460066155
Fax: 0744-2505105
Email: m_tolok@yahoo.com
Web: www.doeacc4u.com

Mahan Computer Services (I) Limited

Established in 1990, the organization is primarily engaged in Education and Training through its own & Franchise centres in the areas of IT Software, Hardware, Networking, Retail Management and English. The institute also provides customised training for corporates.

New Delhi
Tel: 011-25916832-33
Email: info@mahanindia.com
Web: www.mahanindia.com

Enterprise Comm. Solutions

Aware Consultants

We specialize in building and managing Ubuntu/Debian Linux servers and provide good dependable system administration. We install and maintain in-house corporate servers. We also provide dedicated and shared hosting as well as reliable wireless/hybrid networking.

Bangalore
Tel: 080-26724324
Email: sales@aware.co.in
Web: www.aware.co.in

Cynapse India Private Limited

We are the creators of open source product cyn.in. cyn.in is a web 2.0 group collaboration software created by Cynapse, that inter-connects your people with each other and their collective knowledge, seamlessly. It combines the capabilities of collaboration tools like wikis, blogs, file repositories, micro blogs, discussions, audio, videos, and other social applications into a seamless platform. cyn.in helps teams to build collaborative knowledge by sharing and discussing various forms of digital content within a secure, unified application that is accessible using a web based interface or a rich desktop client.

Mumbai
Tel: 022-28445858, 28445629
Email: info@cynapse.com
Web: www.cynapse.com

DeepRoot Linux Pvt Ltd

Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom. We believe in: sharing, compassion and plain action. Backed by full-time hackers. Quick deployment, easy management. Guaranteed.

Bangalore
Tel: 080-40890000
Email: start@deeproot.co.in
Web: www.deeproot.in

ESQUBE Communications Solutions Pvt Ltd

Founders of ESQUBE are faculty at the Indian Institute of Science, Bangalore and carry over eight decades of experience and fundamental knowledge in the field of DSP and Telecommunication. ESQUBE plays a dominant role in the creation of IP in the domain of Sensors, Signals and Systems.

Bangalore
Tel: 080-23517063
Email: info@esqube.com
Web: www.esqube.com

Keen & Able Computers Pvt Ltd

Microsoft Outlook compatible open source Enterprise Groupware Mobile push, Email Syncing of Contacts/Calendar/Tasks with mobiles •Mail Archival •Mail Auditing •Instant Messaging

New Delhi
Tel: 011-30880046, 30880047
Mobile: 09810477448, 09891074905
Email: info@keenable.com
Web: www.keenable.com

netCORE
The Innovation Company

Netcore Solutions Pvt Ltd

No.1 company for providing Linux Based Enterprise Mailing solution with around 1500+ Customer all over India. Key Solutions: •Enterprise Mailing and Collaboration Solution •Hosted Email Security •Mail Archiving Solution •Push Mail on Mobile •Clustering Solution

Mumbai
Tel: 022-66628000
Mobile: 09322985222
Email: kalpit@netcore.co.in
Web: www.netcore.co.in

redhat.

Red Hat India Pvt Ltd

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.

Mumbai
Tel: 022-39878888
Email: marketing-in@redhat.com
Web: www.redhat.in

Hardware & Networking Institute

Xenitis Technolab Pvt Ltd

Xenitis TechnoLab is the first of its kind, state of the art infrastructure, Hardware, Networking and I.T Security training institution headquartered in Kolkata. TechnoLab is the training division of Xenitis group of Companies. It is the proud owner of 'Aamar PC', the most dazzling Desktop brand of Eastern India. These ranges of PC's are sold in the west under the brand name of 'Aamchi PC', in the north as 'Aapna PC' and in the south as 'Namma PC'.

Kolkata
Tel: 033-22893280
Email: solutions@xenitisgroup.com
Web: www.technolabindia.com

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact
Dhiraj (Delhi) **09611206662**, Somaish (Bengalore) **09966076717**

IT Infrastructure Solutions

Absolut Info Systems Pvt Ltd
Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.
New Delhi
Tel: +91 11-26494549
Fax: +91 11-4175 1823
Mobile: +91-9873939960
Email: sales@aisplglobal.com
Web: www.aisplglobal.com

Advent Infotech Pvt Ltd
Advent has an experienced techno-marketing team with several years of experience in Networking & Telecom business, and is already making difference in market place. ADVENT qualifies more as Value Added Networking Solution Company, we offers much to customers than just Routers, Switches, VOIP, Network Management Software, Wireless Solutions, Media Conversion, etc.
New Delhi
Tel: 46760000, 09311166412
Fax: 011-46760099
Email: marketingsupport@adventelectronics.com
Web: www.adventelectronics.com

Asset Infotech Ltd
We are an IT solution and training company with an experience of 14 years, we are ISO 9001: 2000. We are partners for RedHat, Microsoft, Oracle and all Major software companies. We expertise in legal software and solutions.
Dehradun
Tel: 0135-2715965, Mobile: 09412052104
Email: piyush@asset.net.in
Web: www.asset.net.in

BakBone Software Inc.
BakBone Software Inc. delivers complexity-reducing data protection technologies, including award-winning Linux solutions; proven Solaris products; and application-focused Windows offerings that reliably protect MS SQL, Oracle, Exchange, MySQL and other Business critical applications.
New Delhi
Tel: 011-42235156
Email: ashish.gupta@bakbone.com
Web: www.bakbone.com

Clover Infotech Private Limited
Clover Infotech is a leading technology services and solutions provider. Our expertise lies in supporting technology products related to Application, Database, Middleware and Infrastructure. We enable our clients to optimize their Business through a combination of best industry practices, standard processes and customized client engagement models. Our core services include Technology Consulting, Managed Services and Application Development Services.
Mumbai
Tel: 022-2287 0659, Fax: 022-2288 1318
Mobile: +91 99306 48405
Email: business@cloverinfotech.com
Web: www.cloverinfotech.com

DeepRoot Linux Pvt Ltd
Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom. We believe in: sharing, compassion and plain action. Backed by full-time hackers. Quick deployment, easy management. Guaranteed.
Bangalore
Tel: 080-40890000
Email: start@deeproot.co.in
Web: www.deeproot.in

Duckback Information Systems Pvt Ltd
A software house in Eastern India. Business partner of Microsoft, Oracle, IBM, Citrix, Adobe, Redhat, Novell, Symantec, Mcafee, Computer Associates, Veritas, Sonic Wall
Kolkata
Tel: 033-22835069, 9830048632
Fax: 033-22906152
Email: asis@duckback.com
Web: www.duckback.co.in

HBS System Pvt Ltd
System Integrators & Service Provider.Partner of IBM, DELL, HP, Sun, Microsoft, Redhat, Trend Micro, Symentic Partners of SUN for their new startup E-commerce initiative Solution Provider on REDHAT, SOLARIS & JAVA.
New Delhi
Tel: 011 25767117, 25826801/02/03
Fax: 25861428
Email: amittal@hbsindia.com

IgnitionWorks
IgnitionWorks provides customized IT solutions for small/medium businesses by leveraging open source technologies with excellent support & affordable costs. Our expertise include custom software development & infrastructure deployment/support for all industries.
Bangalore
Tel: 080-41743227, Mobile: 9611121432
Email: info@ignitionworks.net
Web: www.ignitionworks.net

INGRES®
Business Open Source

Ingres Corporation
Ingres Corporation is a leading provider of open source database software and support services. Ingres powers customer success by reducing costs through highly innovative products that are hallmarks of an open source deployment and uniquely designed for business critical applications. Ingres supports its customers with a vibrant community and world class support, globally. Based in Redwood City, California, Ingres has major development, sales, and support centers throughout the world, and more than 10 000 customers in the United States and internationally.
New Delhi
Tel: 011-40514199, Fax: +91 22 66459537
Email: sales@ingres.com; info@ingres.com
Web: www.ingres.com

Keen & Able Computers Pvt Ltd
Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.
New Delhi-110019
Tel: 011 30880046, 30880047
Mobile: 09810477448, 09891074905
Email: info@keenable.com
Web: www.keenable.com

LDS Infotech Pvt Ltd
Is the authorised partner for RedHat Linux, Microsoft, Adobe, Symantec, Oracle, IBM, Corel etc. Software Services Offered: •Collaborative Solutions •Network Architecture •Security Solutions •Disaster Recovery •Software Licensing •Antivirus Solutions.
Mumbai
Tel: 022-26849192
Email: sales@ldsinfotech.com
Web: www.ldsinfotech.com

NKIt Solutions & Education
We provide Open Source Solutions, Migration to Linux from any OS and Training in Linux Kernel, Device Driver Development, Linux Administration & Embedded Systems.
Mumbai
Tel: 0184-4032447
Mobile: +91 9896929758, +91 9992211947
Email: info@nkitsolutions.com
Web: www.nkitsolutions.com

Pacer Automation Pvt Ltd
Pacer is leading providers of IT Infrastructure Solutions. We are partners of HP, Redhat, Cisco, Vwmare, Microsoft and Symantec. Our core expertise exists in Consulting, building and Maintaining the Complete IT Infrastructure.
Bangalore
Tel: 080-42823000, Fax: 080-42823073
Email: solutions@pacerautomation.com
Web: www.pacerautomation.com

redhat.

Red Hat India Pvt Ltd
Red Hat is the world's leading open source solutions provider. Red Hat provides high quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.
Mumbai
Tel: 022-39878888
Email: marketing-in@redhat.com
Web: www.redhat.in

Srijan Technologies Pvt Ltd
Srijan is an IT consulting company engaged in designing and building web applications, and IT infrastructure systems using open source software.
New Delhi
Tel: 011-26225926, Fax: 011-41608543
Email: business@srijan.in
Web: www.srijan.in

To advertise in this section, please contact Somaish (Bangalore: 09966075717) Dhiraj (Delhi: 09611206662) on 011-2681-0602 Extn. 222

You **said it**...

tricks (though they are still at it). FOSS is now the race horse chariot in the hands of the Ben Hurs of the world as against the proprietary chariot of the Messiahs

—*J S Karkada, Bangalore, eddykabootar@yahoo.com*

I would have digested Vaibhav's comments easily because all are free to express their views, but the topic "We do not deserve the pride of being open" which my friend selected, forced me to write something politely in reply to his aggressive message. And my message is not for any individual. I want to clear this issue with everyone.

The culture of computing means technology and we the open source users, are far ahead in this field. I have tweaked each and every release of Windows 95, right through to Windows 7 and tens of Linux distros. But there is something in the open world that I don't get in the 'closed' sector. How can we compare Linux to Windows by just considering a few random software? I don't understand why we should even go with GUIs.

A server never demands better GUI effects, all that matters is powerful networking and lawful data management. And if Linux is lagging behind here, then please write and let us know. That is why Linux/UNIX is the first choice of systems administrators.

When developers are busy improving virtualisation why would they divert their efforts to make OpenOffice.org mimic MS Office 2007? The editors like Vim and Emacs are quite challenging to work with, but I love it and that is the real flavour of Linux. Just think of a technology where all are free to develop their programs.

I want to share an experience I had when I installed Fedora 12. My VGA, Ethernet, sound card, Bluetooth webcam, touch pad, card reader, face recognition, and Sify broadband (without adding any package) all worked out of the box. I'm even using Win XP, Win 7, Ubuntu Karmic-Koala, Debian and SUSE (through Xen) and all work fine. So is this an example of 'hard labour'? And then we state: "It's us, the open source community that is lagging behind."

I want to draw your attention to the Grub bootloader, the Linux filesystem, Samba and the emulators. Grub facilitates installing Windows alongside Linux. We can access NTFS drives, files and executables, from within Linux. You might ask why can't an operating system that provides so many facilities offer a better GUI? Well, this is the only thing in Linux for which there has been no demand. Yet, Ubuntu and SUSE are continuously developing better GUIs than Windows—check the sharpness and clarity of their 3D desktops.

The day is not far when all the tech giants will contribute to the open world. Microsoft is doing its work and we are doing ours. There's a lot more to say but I'd like to conclude with: Help FOSS in its work rather than criticise it, because FOSS is *by the users and for the users.*

—*Adarsh Singh, addi_adarsh@yahoo.co.in*

New mag for BSD

As you've done with Linux, I request you to launch a magazine for the BSD community in India. I personally use PC-BSD and Desktop BSD, besides Ubuntu, and I think it would be great to have a BSD magazine.

—*Ashik [mailto:ashiktrissur@aiol.com]*

***ED**: While we appreciate your point, the economics of producing such a niche publication presents a major challenge. However, since LFY is a magazine on open source,*

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact
Dhiraj (Delhi) 09611206662, Somaiah (Bangalore) 09966076717

TechnoInfotech
A company focussed on Enterprise Solution using opensource software.
Key Solutions:
- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

Mumbai
Tel: 022-66338800 Extn: 324
Email: sales@technoinfotech.com
Web: www.technoinfotech.com

Tetra Information Services Pvt Ltd
One of the leading open source providers. Our cost effective business ready solutions caters of all kind of industry verticles.
New Delhi
Tel: 011-46571313 Fax: 011-41620171
Email: sales@tetrain.com
Web: www.tetrain.com

Tux Technologies
Tux Technologies provides consulting and solutions based on Linux and Open Source software. Focus areas include migration, mail servers, virus and spam filtering, clustering firewalls, proxy servers, VPNs, server optimization.
New Delhi
Tel: 011-27348104 Mobile: 09212098104
Email: info@tuxtechnologies.co.in
Web: www.tuxtechnologies.co.in

Veeras Infotek Private Limited
An organization providing solutions in the domains of Infrastructure Integration, Information Integrity, Business Applications and Professional Services.
Chennai
Tel: 044-42210000 Fax: 28144986
Email: info@veeras.com
Web: www.veeras.com

Linux-Based Web-Hosting

Manas Hosting
ManasHosting is a Bangalore-based company that is dedicated in helping small and midsize business companies to reach customers online. We believe that by creating a website, all you have is just web presence; but to get effective traffic on your website it is equally important to have a well designed one. This is why we provide the best of Web Hosting and Web Designing services. Also, our services are backed with exceptionally good quality and low costs.
Bangalore
Tel: 080-42400300
Email: enquiry@manashosting.com
Web: www.manashosting.com

Linux Desktop

Indserve Infotech Pvt Ltd
OpenLx Linux with Kalcutate (Financial Accounting & Inventory on Linux) offers a complete Linux Desktop for SME users. Its affordable (Rs. 500 + tax as special scheme) Friendly (Graphical UserInterface) and Secure (Virus free).
New Delhi
Tel: 011-26014670-71 Fax: 26014673
Email: info@openlx.com
Web: www.openlx.com

Linux Experts

Intaglio Solutions
We are the training and testing partners of RedHat and the first to conduct RHCSS exam in delhi for the first time ever.
New Delhi
Tel: 011-41582917, 45515795
Email: info@intaglio-solutions.com
Web: www.intaglio-solutions.com

Linux Vendor/Distributors

GT Enterprises
Authorized distributors for Red Hat and JBoss range of products. We also represent various OS's Applications and Developer Tools like SUSE, VMWare, Nokia Qt, MySQL, Codeweavers, Ingres, Sybase, Zimbra, Zend-A PHP Company, High Performance Computing Solutions from The Portland Group, Absoft, Pathscale/Qlogic and Intel Compilers Scalix Messaging solution on Linux Platform.
Bangalore
Mobile: +91-9845009938, +91-9343861758
Email: sales@gte-india.com
Web: www.gte-india.com

Taurusoft
Contact us for any Linux Distribution at reasonable rates. Members get additional discounts and Free CD/DVDs with each purchase. Visit our website for product and membership details.
Mumbai
Mobile: 09869459928, 09892697824
Email: taurusoft@gmail.com
Web: www.taurusoft.netfirms.com

Software Subscriptions

Blue Chip Computers
Available Red Hat Enterprise Linux, Suse Linux Enterprise Server / Desktop, JBoss, Oracle, ARCserve Backup, AntiVirus for Linux, Verisign/Thawte/GeoTrust SSL Certificates and many other original software licenses.
Mumbai
Tel: 022-25001812, Mobile: 09821097238
Email: bluechip@vsnl.com
Web: www.bluechip-india.com

Software Development

Carizen Software (P) Ltd
Carizen's flagship product is Rainmail Intranet Server, a complete integrated software product consisting modules like mail sever, proxy server, gateway anti-virus scanner, anti-spam groupware, bandwidth aggregator & manager, firewall, chat server and fax server infrastructure.
Chennai
Tel: 044-24958222, 8228, 9296
Email: info@carizen.com
Web: www.carizen.com

DeepRoot Linux Pvt Ltd
Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom. We believe in: sharing, compassion and plain action. Backed by full time hackers. Quick deployment, easy management. Guaranteed.
Bangalore
Tel: 080-40890000
Email: start@deeproot.co.in
Web: www.deeproot.in

InfoAxon Technologies Ltd
InfoAxon designs, develops and supports enterprise solutions stacks leveraging open standards and open source technologies. InfoAxon's focus areas are Business Intelligence, CRM, Content & Knowledge Management and e-Learning.
Noida
Tel: 0120-4350040, Mobile: 09810425760
Email: sales@infoaxon.com
Web: http://opensource.infoaxon.com

Integra Micro Software Services (P) Ltd
Integra focuses on providing professional services for software development and IP generation to customers. Integra has a major practice in offering Telecom Services and works for Telecom companies, Device Manufacturers, Networking companies, Semiconductor and Application development companies across the globe.
Bangalore
Tel: 080-28856801/06, Fax: 080-28565800
Email: tpvarun@integramicro.com
Web: www.integramicroservices.com

iwebtune.com Pvt Ltd
iwebtune.com is your one-stop, cost with site support organisation. We provide high quality website services and web based software support to any kind of websites, irrespective of the domain or the industry segments.
Bangalore
Tel: 080-4115 8929
Email: santosh@iwebtune.com
Web: www.iwebtune.com

Sarvasy Technologies Pvt Ltd
With ORX Windows Linux Mobile LAMP W3C open, Agile, BDD, TDD development, On-Demand, SaaS, ERP CRM, HRM, Open Source integration, Online project management, VCS (git, svn) automated deployment.
Greater Noida
Mobile: 09711472437, 09810270843
Email: sales@sarvasy.in
Website: www.sarvasy.in

Unistal Systems Pvt Ltd
Unistal is pioneer in Data Recovery Software & Services. Also Unistal is national sales & support partner for BitDefender Antivirus products.
New Delhi
Tel: 011-26288583, Fax: 011-26219396
Email: isales@unistal.com
Web: www.unistal.com

Software and Web Development

Bean eArchitect Integrated Services Pvt Ltd
Application Development, Web Design, SEO, Web Marketing, Web

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact
Dhiraj (Delhi) 09811206562, Somaiah (Bangalore) 09986075717

Development.
Navi Mumbai
Tel: 022-27821617, Mobile: 9820150681
Fax: 022-27821617
Email: infodesk@bonanzhitect.com
Web: www.bonanzhitect.com

Mr Site Takeaway Website Pvt Ltd
Our product is a unique concept in India usingwhich a person without having any technical knowledge can create his website within 1 hour; we also have a Customer Care Center in India for any kind ofafter sales help. We are already selling it world over with over 65,000 copiessold. It comes with FREE Domain Name, Web Hosting and Customer Care Center forFree Support via Phone and Email and features like PayPal Shopping Cart, Guestbook, Photo Gallery, Contact Form, Forums, Blogs and many more. The price ofcomplete package is just Rs 2,999 per year
Patiala
Mobile: 91-9780531855
Email: gurdeep@mrsite.co.in
Web: www.mrsite.co.in

Salah Software
We are specialized in developing custom strategic software solutions using our solid foundation on focused industry domains and technologies. Also providing superior Solution Edge to our Clients to enable them to gain a competitive edge and maximize their Return on Investments (ROI).
New Delhi
Tel: 011-41648668, 66091560
Email: ceo@salahsoftware.com
Web: www.salahsoftware.com

Thin Client Solutions

Digital Waves
The 'System Integration' business unit offers end-to-end Solutions on Desktops, Servers, Workstations, HPC Clusters, Render Farms, Networking, Security/Surveillance & Enterprise Storage. With our own POWER-X branded range of Products, we offer complete Solutions for Animation, HPC Clusters, Storage & Thin-Client Computing
Mobile: 09880715253
Email: rangesh@digitalwaves.in
Web: www.digitalwaves.in

Enjoy Network Solutions
Gujarat based ThinClient Solution Provider. Providing Small Size ThinClient PCs & a Full Featured ThinClient OS to perfectly suite needs of different working environment. Active Dealer Channel all over India
Gujarat
Tel: 0265-3205400, 3201732, 3251732,
Mobile: 09377107020, 09898007850
Email: info@enjoyworld.com
Web: www.enjoyworld.com

Training for Corporate

Bascom Bridge
Bascom Bridge is Red Hat Certified partner for Enterprise Linux 5 and also providing training to the individuals and corporate on other open source technologies like PHP, MySQL etc.
Ahmedabad
Tel: 079-27545455—66
Fax: 079-27545488
Email: info@bascombridge.com
Web: www.bascombridge.com

Brainnet
Kolkata
Tel: 033-40076450
Email: brainnet@brainware-india.com
Web: www.brainware-india.com

Centre for Excellence in Telecom Technology and Management (CETTM), MTNL
MTNL's Centre for Excellence in Telecom Technology and Management (CETTM) is a state of the art facility to impart Technical, Managerial and corporate training to Telecom; Management personnel. CETTM has AC lecture halls, computer Labs and residential facility
Mumbai
Tel: 022-25714500, 25714586, 25714585, 25714580
Fax: 022-25706700
Email: contact@cettm.mtnl.in
Web: http://cettm.mtnl.in/infra

Complete Open Source Solutions
RHCT, RHCE and RHCSS training.
Hyderabad
Tel: 040-66773365, 9849742065
Email: nayak.sujeet@gmail.com
Web: www.cossindia.com

ElectroMech
Redhat Linux and open source solution, RHCE, RHCSS training and exam center,Ahmedabad and Vadodara
Ahmedabad
Tel: 079-40027898
Email: electromech@electromech.info
Web: www.electromech.info

Focuz Infotech
Focuz Infotech Advanced Education is the quality symbol of high and Advanced Technology Education in the state. We are providing excellent services on Linux Technology Training, Certifications and live projects to students and corporates, since 2000.
Cochin
Tel: 0484-2395324
Email: enquiry@focuzinfotech.com
Web: www.focuzinfotech.com

G-TEC Computer Education
ISO 9001:2000 certified IT Company, International Testing Centre, Specialised in Multimedia & Animation, conduct MCP, MCSE 2000, MCDBA and MCSA certificates, CCNA, CCNP, the Only authorized centre by INTERNATIONAL AND EUROPEAN COMPUTER UNION to conduct ICDL, Adobe Certifications, training on Web Designing, Tally, Spoken English, Conducts Corporate and Institutional training, International certifications issued.
Bangalore
Tel: 080-42867300
Email: gtec.indiranagar@gmail.com

Gujarat Infotech Ltd
GIL is a IT company and 17 years of experience in computer training field. We have experience and certified faculty for the open Source courses like Redhat, Ubantoo,and PHP, Mysql
Ahmedabad
Tel: 079-27452276, Fax: 27414250
Email: info@gujaratinfotech.com
Web: www.gujaratinfotech.com

Lynus Academy Pvt Ltd
India's premier Linux and OSS training institute.
Chennai
Tel: 044-42171278, 9840880558
Email: contactus@lynusacademy.com
Web: www.lynusacademy.com

Linux Learning Centre Private Limited
Pioneers in training on Linux technologies.
Bangalore
Tel: 080-22428538, 26600839
Email: info@linuxlearningcentre.com
Web: www.linuxlearningcentre.com

Maze Net Solutions (P) Ltd
Maze Net Solution (P) Ltd, is a pioneer in providing solutions through on time, quality deliverables in the fields of BPO, Software and Networking while providing outstanding training to aspiring IT Professionals and Call Center Executives. Backed by a team of professional workforce and global alliances, our prime objective is to offer the best blend of technologies in the spheres of Information Technology (IT) and Information Technology Enabled Services (ITES).
Chennai
Tel: 044-45982535
Email: info@mazenetsolution.com
Web: www.mazenetsolution.com

Netweb Technologies
Simplified and scalable storage solutions.
Bangalore
Tel: 080-41144595, 32718516
Email: info@netwebindia.com
Web: www.netwebindia.com

New Horizons India Ltd
New Horizons India Ltd, a joint venture of New Horizons Worldwide, Inc. (NASDAQ: NEWH) and the Shriram group, is an Indian company operational since 2002 with a global foot print engaged in the business of knowledge delivery through acquiring, creating, developing, managing, lending and licensing knowledge in the areas of IT, Applied Learning, Technology Services and Supplementary Education. The company has pan India presence with 15 offices and employs 750 people.
New Delhi
Tel: 011-45612400
Email: info@nhindia.com
Web: www.nhindia.com

Network NUTS
India's only Networking Institute by Corporate Trainers. Providing Corporate and Open classes for RHCE / RHCSS training and certification. Conducted 250+ Red Hat exams with 95% result in last 8 months. The BEST in APAC.
New Delhi
Tel: 46526980-2
Mobile: 09310024503, 09312411592
Email: info@networknuts.net
Web: www.networknuts.net

STG International Ltd
An IT Training and Solution Company Over an experience of 14years.We are ISO 9001:2000

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact
Dhiraj (Delhi) 09611206582, Somaiah (Bangalore) 09966076717

Certified/Authorised Training Partners of Red Hat & IBM-CEIS. We cover all Software Trainings
New Delhi
Tel: 011-42560441-42, Mobile: 09873108801
Email: nakul@sig.in
Web: www.sigonline.com
www.sigglobal.com

TNS Institute of Information Technology Pvt Ltd

Join RedHat training and get 100% job gaurantee. World's most respected Linux certification. After RedHat training, you are ready to join as a Linux Administrator or Network Engineer.
New Delhi
Tel: 011 3085100, Fax: 30851133
Email: nks@tiit.co.in
Web: www.tiit.co.in

Webel Informatics Ltd

Webel Informatics Ltd (WIL), a Government of West Bengal Undertaking. WIL is Red Hat Training Partner and CISCO Regional Networking Academy. WIL conducts RHCE, RHCSS, CCNA, Hardware and Software courses.
Kolkata
Tel: 033-22833568, Mobile: 09433111110
Email: enquiry@webelinformatics.com
Web: www.webelinformatics.com

Training for Professionals

AEM

AEM is the Best Certified Redhat Training Redhat in Eastern India since last 3 years. AEM conducted more than 500 RHCE exams with 95-100% pass rate. Other courses— RHCSS,SCNA,MCSE CCNA.
Kolkata
Tel: 033-25488736, Mobile: 09830075018
Email: sinhatuhin1@gmail.com
Web: www.aemk.org

Agam Institute of Technology

In Agam Institute of Technology, we provide hardware and networking training since last 10 years. We specialise in open source operating systems like Red Hat Linux since we are their preferred training partners.
Dehradun
Tel: 0135-2673712, Mobile: 09760099050
Web: www.agamtecindia.com

Amrita Institute of Computer Technology

Amrita Technologies provides an extensive training in high end certification programs and Networking Solutions like Redhat Linux, Redhat Security Services, Cisco, Sun Solaris, Cyber Security Program IBM AIX and so on with a strong focus on quality standards and proven technology processes with most profound principles of Love and Selfless Service.
Mobile: 09393733174
Email: aid.hyb@amrita.ac.in
Web: www.amritahyd.org

Centre For Industrial Research and Staff Performance

A Unique Institute catering to the need for industries as well as Students for trainings on IT, CISCO certification, PLC, VLSI, ACAD, Pneumatics, Behavior Science and Handicraft.
Bhopal
Tel: 0755-2661412, 2661559
Fax: 0755-4278622
Email: crisp@crispindia.com
Web: www.crispindia.com

Center for Open Source Development And Research

Linux, open source & embedded system training institute and development. All trainings provided by experienced experts & administrators only. Quality training (corporate and individual). We expertise in open source solution.Our cost effective business ready solutions caters of all kind of industry verticals.
New Delhi
Mobile: 09312506496
Email: info@cfosdr.com
Web: www.cfosdr.com

Cisconet Infotech (P) Ltd

Authorised Red Hat Study cum Exam Centre. Courses Offered: RHCE, RHCSS, CCNA, MCSE
Kolkata
Tel: 033-25395508, Mobile: 09831705913
Email: info@cisconetinfo.com
Web: www.cisconetinfo.com

CMS Computer Institute

Red Hat Training partner with 3 Red Hat Certified Faculties, Cisco Certified (CCNP) Faculty, 3 Microsoft Certified Faculties having state Of The Art IT Infrastructure Flexible Batch Timings Available. Leading Networking Institute in Marathwada
Aurangabad
Tel: 0240-3299509, 6621775
Email: aurangabad@cmsinfotech.co.in
Web: www.cmsaurangabad.com

Cyber Max Technologies

OSS Solution Provider, Red Hat Training Partners, Oracle,Web, Thin Clients, Networking and Security Consultancy. Also available CCNA and Oracle Training on Linux. Also available Laptops & PCs
Bikaner
Tel: 0151-2202105, Mobile: 09928173269
Email: cmtech.bikaner@gmail.com
or gupta.vaibhav@gmail.com

Disha Institute

A franchisee of Unisoft Technologies Providing IT Training & Computer Hardware & Networking
Dehradun
Tel: 3208054, 09897168902
Email: thedishainstitute@gmail.com
and www.unisofttechnologies.com

EON Infotech Limited (TECHNOSchool)

TechnoSchool is the most happening Training Centre for Red Hat (Linux, Open Source) in the Northern Region. We are fully aware of the Industry's requirement as our Consultants are from Linux Industry. We are committed to make you a total industry ready individual so that your dreams of a professional career are fulfilled.
Chandigarh
Tel: 0172-5067566-67, 2609849
Fax: 0172-2615465
Email: info@technoschool.net
and: http://technoschool.net

GT Computer Hardware Engineering College (P) Ltd

Imparting training on Computer Hardware Networking, Mobile Phone Maintenance & International Certifications
Jaipur
Tel: 0141-3213378
Email: franchise_gt@gteducation.net
and: www.gteducation.net

HCL Centre Development Centre Bhopal

As the fountainhead of the most significant pursuit of human mind (IT), HCL strongly believes, "Only a Leader can transform you into a Leader". HCL CDC is a formalization of this experience and credo which has been perfected over three decades.
Bhopal
Tel: 0755-4094852
Email: bhopal@hclcdc.in
Web: www.hclcdc.in

FOSS Yellow Pages

INZTRIX E Technologies Pvt Ltd

No. 1 Training provider in this region.
meerut
Tel: 0161-4000111, 4000222
Mobile: 09927666664
Email: jai@inirtx.com
Web: www.inirtx.com

Indian Institute of Job Oriented Training Centre

Ahmedabad
Tel: 079-40072244—2255—2266
Mobile: 09898749595
Email: info@iijt.net
Web: www.iijt.net

Institute of Advance Network Technology (IANT)

•Hardware Engg.•Networking
•Software Engg. •Multimedia Training.
Ahmedabad
Tel: 079-32516577, 26607739
Fax: 079-26617738
Email: contact @iantindia.com
and www.iantindia.com

IPCC

Bridging Gap with professionals.
Lucknow
Tel: 0522-3013496
Email: ipcclko@yahoo.co.in
Web: www.ipcc.co.in

Koenig Solutions (P) Ltd

A reputed training provider in India. Authorised training partner of Red Hat, Novell and Linux Professional Institute. Offering training for RHCE, RHCSS, CLP, CLE, LPI - 1 & 2.
New Delhi
Mobile: 09911713143, Fax: 011-25639990
Email: info@koenig-solutions.com
Web: www.koenig-solutions.com

NACS/CIT

We are Providing Training of LINUX to Professional & Corporate.
meerut
Tel: 0121-2420587, Mobile: 9997526668
Email: info@nacsglobal.com
and www.nacsglobal.com

NACS Infosystems (P) Ltd

NACS is a organization which is providing training for all International certification, and also NACS is the authorized Training Partner of Redhat and also having testing centre of THOMSON PROMETRIC and PEARSON VUE.
Meerut
Tel: 0121-2767756, Fax: 0121-4006551
Mobile: 09897796603
Email: info@nacsglobal.com

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact
Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

mohit@nacoglobal.com.
Web: www.nacoglobal.com

Netdlox Computing Systems
We are one-of-a-kind center for excellence and finishing school focusing on ground breaking technology development around distributed systems, networks, storage networks, virtualisation and fundamental algorithms optimised for various appliance.
Bangalore
Tel: 080-28540700
Mobile: 09743045886
Email: info@netdlox.com

Netdlox Technologies
Training Partner of RedHat,Cisco
Chandigarh
Tel: 0172-2608851, 3916555
Email: mail.netmax@gmail.com
Web: www.netmaxtech.com

Netxpert Institute of Advance Networking
Netxpert Mobile is a Leading organization to provide Open Source training on RedHat Linux RHCT and RHCE Training with 30Hrs. extra exam preparation module.
Noida
Tel: 0120-4346847 Mobile: 09268829812
Email: anoop.vyas@netxprtindia.com
Web: www.netxprtindia.com

Netzone Infotech Services Pvt Ltd
Special batches for MCSE, CCNA and RHCE on RHEL 5 with exam prep module on fully equipped labs including IBM servers, 20+ routers and switches etc. Weekend batches are also available.
New Delhi
Tel: 011-46015674 Mobile: 9212114211
Email: info@netzoneindia.net

Source IT Solutions
We offer end to end services and support to implement and manage your IT Infrastructure needs. We also offer Consulting services and Training in Advanced Linux Administration
Chennai
Mobile: 09790804000
Email: enquiry@sourceit.in
Web: www.sourceit.in

Plexus Software Security Systems Pvt Ltd
Plexus, incorporated in January 2003 is successfully emerged as one of the best IT Company for Networking, Messaging & Security Solutions and Security Training. Networking, Messaging & Security solutions is coupled with the expertise of its training, this has put Plexus in the unique position of deriving synergies between Networking, Messaging & Security Solutions and IT Training.
Chennai
Tel: 044-2433 7355
Email: training@plexus.co.in
Web: www.plexus.co.in

Professional Group of Education
RHCE & RHCSS Certifications
Jabalpur
Tel: 0761-4039376,
Mobile: 09425152831
Email: naidu.vikas@gmail.com

Q-SOFT Systems & Solutions Pvt Ltd
Q-SOFT is in a unique position for providing technical training required to become a Linux Administration under one roof. Since inception, the commitment of Q-SOFT towards training is outstanding. We Train on Sun Solaris, Suse Linux & Redhat Linux.
Bangalore
Tel: 080-26639207/26544135, 22440507
Mobile: +91 9945 282834
Email: counselors@qsoftindia.com
Web: www.qsoftindia.com

Software Technology Network
STN is one of the most acknowledged name in Software Development and Training. Apart from providing Software Solutions to various companies, STN is also involved in imparting High-end project based training to students of MCA and B.Tech etc. of various institutes.
Chandigarh
Tel: 0172-5086829
Email: stn2001@rediffmail.com
Web: stntechnologies.com

South Delhi Computer Centre
SDCC is for providing technical training courses (software, hardware, networking, graphics) with career courses like DOEACC "O" and "A" Level and B.Sc(IT),M.Sc(IT),M.Tech(IT) from KARNATAKA STATE OPEN UNIVERSITY.
New Delhi
Tel: 011-26183327, Fax: 011-26143642
Email: southdelhicomputercentre@gmail.com, southdelhicomputercentre@hotmail.com.
Web: www.itwhizkid.com
www.itwhizkid.org

Sysfema Quest
Making Tomorrow's professionals TODAY
Bangalore
Tel: 080-41301814
Email: directors@sysfemaquest.com
Web: www.sysfemaquest.com

Trimax FuturePerfect
A Div. of Trimax IT Infrastructure and Services Limited. Redhat RHCE, RHCT Training & Exam Center MCTS, MCITP, MCSE 03, CCNA, CCNP, Prometric Center.
Mumbai
Tel: 022-40681313, Mobile: 09987709955
Fax: 022-40661001
Email: futureperfect@trimax.in
Web: www.trimax.in

Vibrant e Technologies Ltd
Vibrant e Technologies Ltd. Is a authorised Red Hat Test and Testing Centre, has won the prestigious award " REDHAT BEST CERTIFIED TRAINING PARTNER 2007-2008" for Western region. Vibrant offers courses for RHCE & RHCSS etc.
Mumbai
Tel: 022-26285066/6701
Email: vibrant@vsnl.net
Web: www.vibrantcomputers.com

Ultramax Infonet Technologies Pvt Ltd
Training in IT related courses adn authorised testing center of Prometric, Vue and Red Hat.
Mumbai
Tel: 022-67669217
Email: unmesh.raote@ultramaxit.com
Web: www.ultramaxit.com

Yash Infotech
Authorized Training & Exam Center. Best Performing Center in Lucknow for RH Training and Examinations. LINUX & Open Source training institute for IT professionals & Corporate Offering Quality Training for RHCE, RHCSS, PHP, Shell Script, Virtualization and Troubleshooting Techniques & Tools.
Lucknow
Tel: 0522-4043386, Fax: 0522-4043386
Email: yashinfotech.lko@gmail.com

Web Hosting

IDS Logic Pvt Ltd
IDS Logic is a leading innovative IT Solutions company provides Linux Web Hosting, Windows Web Hosting Application Hosting, Reseller Hosting Dedicated Server, Virtual Private Server Server Management & Monitoring, Domain Name Registration, SSL Certificate, Corporate Mailing Solutions, Web Security Solutions.
Noida
Tel: 0120-4229865, Fax: 0120-4235695
Email: info@idslogic.com
Web: www.idswebhosting.com

Perfect Innovation
Web Hosting Spider
- Web Hosting • Web Design
- Web Application Development
- SMS Hosting • Corporate Hosting
- Dedicated Servers

Puducherry
Tel: 0413-3202726, 3246999
Email: admin@webhostingspider.com
Web: www.webhostingspider.com

YOU CAN CHOOSE ANY OF THE FOLLOWING OPTIONS

Pay for 6 months Get for 1 year!

Pay for 1.5 years Get for 3 years!

Pay for 2.5 years Get for 5 years!

EFY PUBLICATIONS

A magazine for the electronics fraternity: professionals, businessmen and hobbyists

India's first Electronics Sourcing magazine

Asia's first magazine on Linux and Open Source

A monthly on business & economic affairs with regular sections on market surveys

MEGA SUBSCRIPTION OFFER

MAGAZINE	☑ Mark Required	Unit Price (Rs.)	Pay for 6 months Get for 1 Year	Pay for 1.5 Years Get for 3 Years	Pay for 2.5 Years Get for 5 Years
Electronics For You (with CD)	☐	60/-	☐ Pay Rs. 360/-	☐ Pay Rs. 1080/-	☐ Pay Rs. 1800/-
Electronics Bazaar	☐	50/-	☐ Pay Rs. 300/-	☐ Pay Rs. 900/-	☐ Pay Rs. 1500/-
Linux For You (with CD & DVD)	☐	100/-	☐ Pay Rs. 599/-	☐ Pay Rs. 1800/-	☐ Pay Rs. 3000/-
Facts For You	☐	100/-	☐ Pay Rs. 599/-	☐ Pay Rs. 1795/-	☐ Pay Rs. 2995/-

• These rates are applicable for new subscribers, as well as renewal by existing subscribers. • The rates are valid for subscribers within India only • Allow 4-6 weeks for despatch of your subscription. • Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only.

Hurry!

Offer Valid till 31st March 2010

Subscribe Now!

ORDER FORM

Name ______ Company Name ______ Designation ______

Mailing address ______ Dist. ______

City ______ Pin Code ______ State ______ Phone ______

Email ______ Subscription No. (for existing subscribers only) ______

I would like to subscribe to the above (✓)marked magazine(s), for the tenure ______ Please find enclosed a sum of Rs ______ by DD/MO/crossed cheque*bearing the No. ______ dt. ______ in favour of EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs 50 on non-metro cheque)

or Please charge my credit card ☐ VISA ☐ MasterCard Please charge Rs. ______

against my credit card No. └┘└┘└┘└┘└┘└┘└┘└┘└┘└┘└┘└┘└┘└┘└┘└┘

Date of Birth ___/___/______ (dd/mm/yy) Card Expiry Date ______/______ (mm/yy)

Signature (as on the card)

Send this filled-in form or its photocopy to
EFY Enterprises Pvt Ltd, D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020
Ph: 011-26810601-03; Fax: 011-26817563; E-mail: info@efyindia.com website: www.efyindia.com

R N I No DELENG/2003/9889, Mailed on 1/2nd of the month
Published on 28th of the previous month
Postal Regd. No. DL(S)-01/3001/2009-11
Licenced to Post without Prepayment Licence No. U(SE)-04/2009-11

You **said it**…

we do cover BSD too. It's just that very few people write about it. If you (along with friends you know in the BSD community) can contribute articles on the subject, we could have a regular series on BSD in LFY and on LinuxForU.com too. By the way the DVD this month contains the ISO image of the latest Free BSD 8.0 release. Hope you'll burn it to a DVD to install it, and then, perhaps, send us feedback on what you think about the new features from the land of BSD :)

Maemo coverage

I recently went through the Nokia Maemo website (*maemo.nokia.com*). Can you do a review of Maemo (and optionally the N900)? It appears to me to be one of the best demonstrations of the capabilities of Linux and other FOSS products such as Mozilla based browsers, OpenGL, etc. Also, a primer on developing apps for Maemo would be a bonus.

—Saurav Sengupta, Chandigarh, saurav.public@gmail.com

ED: *Well, you've seen the February 2010 cover already. Hope you like the coverage given to Maemo.*

Online CAT

Recently the online CAT was in the news quite a lot. We use Moodle as a learning management system and specifically to conduct online exams. Initially we too faced problems similar to what CAT faced. The cause of the problems was the simultaneous access by a large number of users distributed across the country.

We have an LVS with a back up for load balancing. It connects to two Web servers (Apache), which in turn connect to a DB server (MySQL). We also have a storage server in the system. All of them are interconnected through a high speed LAN. Even though our hardware can handle a large number of simultaneous accesses, Moodle itself has limitations. Bandwidth constraints add to the problem. Any amount of server tuning and application tuning does not solve the problem.

I believe that the CAT online exam faced similar problems. We solved our problems by scheduling the exams such that not more than about 700 people took it at one time. Where the examinees are located far away we set up local servers and synchronised with the master server manually. Though the last step can be automated, I think some amount of manual intervention is healthy. We also created a large question bank where questions are picked up randomly, so that different groups got different question papers.

—Lalitha K Prasad, Head, Corporate Learning Centre, Tata Consultancy Services

ED. *Thanks—this is very interesting feedback.*

ERRATA January 2010

The word 'bootstrapping' that appeared on the cover of the last issue and in the section highlighting the contents on page 4 was incorrectly spelt as 'bootstraping'. We regret this error and promise to take steps to avoid such silly mistakes.

Please send your comments or suggestions to:
The Editor
LINUX FOR YOU Magazine
D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020
Phone: 011 26810601/02/03, **Fax**: 011 26817563
Email: lfyedit@efyindia.com, **Website**: www.openITis.com

Q&A SECTION

Q I have a few queries regarding my first installation of Fedora 12 on my PC. Currently, I am using Windows XP SP2, and my motherboard is MSI-7191(v1.x), which supports an AMD K8 ATHLON 64FX processor. Does Fedora require the installation of any motherboard driver CD? If it does not, then how will I enable my sound, networking, and enhance my display features? Also, how do I connect to the Internet through my cellphone using PC Suite? It's an EXE file, and I've heard Linux doesn't support this format. Can you tell me where I can get a PDF or Word file that tells me about the Fedora 12 installation process and things I need to do, post installation?

—Dhananjay Sharma, dhananjaykrm@gmail.com

As far as we know, most of the features work out of the box for MSI-7191 (but you need to install and check this, as we have not tested it). This means that you do not need any additional drivers for your motherboard. Your LAN and sound should also work out of the box.

To connect to the Internet using your mobile, you first need to check that your Linux installation recognises your phone modem. If it does, then you can use any dialler software to connect. Try the NetworkManager program—it's quite easy to use, and you should be able to go online with it. Google for a how-to; there's plenty of documentation on it. You can even join the Fedora Forum at *http://fedoraforum.org* and search for solutions to the issues that persist on your system. Only when you can't find a solution, post your query to the forum.

As for a Fedora 12 installation guide, you'll find it at. *http://docs.fedoraproject.org/install-guide/f12/en-US/html/*

Q I have installed openSUSE 11.1, and am happy with its performance. I have another machine on my network with Windows XP installed on it. I was trying to share my folder on my Linux machine so that it can be accessed on the Windows PC. I tried to share the folder with the option available on the Nautilus file manager (when you right click on the folder and select *Properties*). However, upon configuring the share this way, it shows an error message. Please suggest a solution for this.

—Tanya Verma, New Delhi

First of all, check that you have the following packages installed on your system: *samba, samba-client and nautilus-share*. If not, then install them first. Now open the */etc/samba/smb.conf* file in a text editor and add the following lines under the [global] section, if they aren't there:

```
usershare allow guests = Yes
usershare max shares = 100
usershare owner only = False
```

Remember, you'll only be able to edit this file as the root user; do not forget to take a back-up copy of the *smb.conf* file before you make any changes to it. Now check whether you have the */var/lib/samba/usershares/* directory as configurations get stored here. If you don't, run the following command as the root:

```
mkdir /var/lib/samba/usershares
chgrp users /var/lib/samba/usershares
chmod 1770 /var/lib/samba/usershares
```

The last step is to restart your Samba server, as follows:

```
/etc/init.d/smb restart
```

Now, try to share the folder using Nautilus—it should work.

Q I have got a live CD of openSUSE 11.2, but I don't know how to install it, keeping Windows as it is. Can you please help install it alongside Windows?

—Vikas, clevervikas@gmail.com

To install openSUSE, all you need to do is boot your system using the Live CD and then click on the *Live Installer* icon on your live desktop to start the installation. Follow the on-screen instructions. On the partitioning screen, you can either customise the partitioning, or you can accept the default partitioning.

I would suggest you carefully review the partitioning section to see that openSUSE should install in an empty partition. [If you're unsure—you don't want to wipe out your data by selecting a wrong partition to install openSUSE—take the help of a friend, Google, or let us know about your partition table.] The Installer will automatically select your Windows partition and set it up for a dual boot system. Now follow the on-screen instructions to complete the installation. After it's over, you should have an option to select between Windows and openSUSE when you boot your PC.

HACKING

FOSS Bytes

Google starts selling Nexus One, not in India though

Taking on rival Apple's iPhone, Internet search giant Google has launched its own smartphone which according to the company "belongs in the emerging class of devices" that it calls superphones. The Google Nexus One combines the latest in hardware from HTC Corporation with the newest Android software (Android 2.1).

Manufactured by HTC of Taiwan, the Nexus One features noise suppression from Audience, Inc; a 9.4 cm (3.7inch) AMOLED 480 x 800 pixel WVGA display (a larger screen than Apple's iPhone) and a 1GHz Qualcomm Snapdragon chipset for blazing speeds. Weighing only 130 grams, the phone also includes a 5 MP camera with auto focus, a flash and geo tagging; onboard memory of 512 MB Flash, 512MB RAM and a 4GB MicroSD Card (expandable to 32 GB); 3.5-mm stereo headphone jack and 1400 mAh battery. The trackball doubles as a notification tool, which contains a multi-colour LED for indicating incoming calls and other notifications.

Running on Android 2.1, the superphone includes a voice-enabled keyboard so you can speak into any text field, fun Live Wallpapers, a 3D photo gallery for richer media experiences and lots more. Of course, it also comes with a host of Google applications, including Gmail, Google Voice and Google Maps Navigation.

The Nexus One is initially available from the Google Web store in the US—$529 for an unlocked phone, or starting at $179 with a two-year contract from T-Mobile USA. In the near future, Verizon Wireless in the US and Vodafone in Europe plan to offer services to customers in their respective geographies.

MSI to ship SUSE Moblin-powered netbook

MSI and Novell have announced that SUSE Moblin preloaded on the MSI U135 netbook will be launched some time in February 2010. Following the recent release of Moblin version 2.1, this marks the first OEM to sell a fully-supported netbook, with the Intel Atom processor running Moblin based technology. The netbook offers a rich Internet experience with integrated e-mail, multimedia and Web browsing capabilities. The U135 netbook claims to offer nearly all-day computing capabilities with long battery life and low power consumption. For more information on SUSE Moblin, visit *www.novell.com/moblin*.

World's first Snapdragon smartbook to run Android

Qualcomm has announced it is working with HP to design an Android-based smartbook device. The design utilises Qualcomm's Snapdragon QSD8250 chipset platform with an integrated Scorpion CPU, delivering speeds up to 1 GHz. The device demonstrates how long battery life, 3G and Wi-Fi connectivity, a highly intuitive user interface and other features can meet the consumer's need for a mobile yet constantly connected lifestyle. HP demonstrated the device at the 2010 Consumer Electronics Show in Las Vegas.

Qualcomm's Snapdragon platform claims to offer an unprecedented combination of gigahertz mobile processing performance, optimised power consumption, ubiquitous connectivity and powerful multimedia, all on a single chip.

Linux Foundation starts online jobs board

The demand for Linux professionals has grown multifold since 2005. To cater to the rising demand for Linux related jobs, the Linux Foundation (LF) has come up with its new Linux Jobs Board (*http://jobs.linux.com*). This will provide employers and job seekers with an online forum in which anyone can find the best and brightest Linux talent or the ideal job opportunity.

Employers interested in posting a job opening have two options: they can post their openings on Linux.com for $100 for 15 days or $200 for 30 days and reach millions of Linux-focused professionals, or pay additional fees to post jobs on the JobThread Network on Linux.com to reach 50 niche publishing sites with a combined 9.8 million visitors every month, reports LinuxDevices.com.

EU set to clear Oracle-Sun deal

MySQL founder Michael Widenius is taking his vocal campaign against Oracle's planned takeover of Sun Microsystems to China and Russia because the European Commission appears set to clear the deal, reports Reuters. The MySQL open source database is owned, developed and supported by Sun. The Commission was initially concerned that Oracle's takeover of the MySQL database could hurt competition in that market. However, Oracle recently said it would continue to enhance MySQL and would release future versions under an open source licence. The EC is likely to clear the deal by January 27, 2010.

FOSS Bytes

Freescale releases reference design for under-$200 smartbooks

Freescale has unveiled what could become the future of the smartbook category, showcasing a tablet reference design with a 17.7 cm (7 inch) touch screen with up to four times the viewing area of a typical smartphone. It is based on a form factor that is approximately one third the size and volume of today's notebooks. Details of the reference design are available at *www.freescale.com/smartbook* and it is expected to be available for evaluation beginning February 2010.

The solution is based on the Freescale i.MX515 processor that incorporates the ARM Cortex-A8 technology. It also includes Freescale MC13892 power management IC, SGTL5000 audio codec and MMA8450Q 3-axis accelerometer. The Freescale tablet includes Wi-Fi and Bluetooth wireless connectivity and features a 3D desktop framework with touchscreen/QWERTY keyboard support. Options for a 3G modem and RF4CE protocol are available.

The solution will enable a second generation of smartbook products priced below $200 and featuring form factors that fully leverage the power, performance and functionality of advanced ARM processor technology. It is designed to provide instant-on functionality, persistent connectivity and all-day battery life. The tablet was demonstrated at the 2010 Consumer Electronics Show in Las Vegas in January running both the Android and other Linux operating systems.

Get an Ubuntu netbook for Rs 13,999

Simmtronics Semiconductors has signed a global deal with IBM as part of which IBM has authorised Simmtronics to pre-load its *IBM Client for Smart Work* software solution in its Simmbook that runs Ubuntu. This includes Lotus Symphony office productivity software. Under the agreement, both the companies will execute a joint marketing plan and work closely on the sales opportunities for Simmbook with *IBM Client for Smart Work* on Ubuntu, globally.

Priced at Rs 13,999, the Simmbook is now available in India through the company's 16 direct sales offices. The netbook is powered by a 1.6 GHz Intel Atom processor and Intel Generation 3.5 Integrated GFX Core Graphic VGA. It comes in a black, blue and red colour chassis, with a 25.9 cm (10.2 inch) LCD screen and 1024 X 600 pixel resolution. Other device specs include: 160 GB SATA HDD, 1GB DDR2 667MHz RAM, a 1.3 MP camera, a 3-in-1 MMC/SD/MS/MS Pro card reader and three USB 2.0 ports. It weighs about 1.2 kg with a three-cell battery pack. For more details, visit *www.simmtronics.com/simmbook.php*

French govt to citizens: Stop using Internet Explorer

The French government has issued an advisory to computer users, recommending that they switch to a different Web browser, such as Firefox or Google Chrome. France's move follows that of the German government, which recently discovered that Internet Explorer contained a dangerous security flaw that could be exploited by cyber criminals.

Microsoft admitted that its Internet Explorer 6 browser was the weak link in recent attacks that cracked into e-mail accounts of human rights activists in China. But it said that the German government had gone too far in its over-reaction, and that general users were not at risk. But that did not stop the French government issuing a bulletin through Certa, a government agency responsible for managing and protecting against cyber threats. The bulletin said that all versions of Internet Explorer were affected by the security flaw and suggested people use an alternative browser.

New Pardus out!

Pardus Linux 2009.1 comes with the latest stable KDE release, enhanced hardware support, and a bigger software archive with up-to-date packages such as KDE 4.3.4, Linux kernel 2.6.31.11, OpenOffice.org 3.1.1, Mozilla Firefox 3.5.7, the GIMP 2.6.8, X.Org Server 1.6.5, Python 2.6.4 and many more, in just one CD. With the 2009.1 series, the Pardus ISO files are generated as hybrid images, which can be burned to a CD, DVD or dumped to a hard disk like media, such as a USB stick.

FOSS Bytes

Now, an Android module targets everyday appliances

San Francisco-based Touch Revolution has unveiled its new drop-in NIM1000 module, which is designed to help OEMs integrate dynamic touch-screen interfaces in a broad range of appliances and devices, from microwave ovens and washing machines, to in-flight entertainment centres and advanced medical devices. The NIM1000 module includes a projective capacitive touch-screen and features the Android application framework, which makes it easy for manufacturers to add a myriad additional applications and services to their products. To demonstrate how the drop-in NIM1000 module might transform next-generation appliances and devices, Touch Revolution showcased several potential products at CES 2010 in Las Vegas.

- A washing machine and dryer with simple, slider touch-screen controls to set wash and dry cycles and view progress. Online stain guides provide information on removing the toughest stains and those cryptic label wash symbols are decoded in a few touches.
- A microwave oven that doubles as a kitchen command centre with Android running all kinds of useful applications, widgets for news and weather, a digital picture frame, Internet radio, recipes with video instructions, a Web browser and timers for each burner and the oven.

Yes, there was more, and if you let your imagination run wild, you should be able to come up with a few yourself.

Moodbar and other custom labels back in Amarok 2.2.2

What's the difference between Amarok 2.2.2 and 2.2.1? If "Bug fix release," is your answer, then you're quite right. However, Amarok 2.2.2 also introduces some of the features that were once available in the v1.4.x series of the media player but were absent since v2.x made its debut. Here's quoting the announcement at the project home page: "...we focused on improvements and bug fixes, but let us remind you of the new features in this release first. The Custom Labels are back and allow you to add your personal touch to your music. Another new addition that will add quite some colour to your music experience: the moodbar makes its entry in the 2.2 series so you can change the progress bar to fit with your current mood, be this normal, angry, frozen or just happy :)"

That's not all! There are a number of small, but useful improvements and feature additions that makes this version worth an upgrade.

VMware to acquire Zimbra

VMware has announced that it has entered into a definitive agreement to acquire Zimbra, from Yahoo! The virtualisation giant claims this acquisition will further its mission of taking complexity out of the data centre, desktop, application development and core IT services, and delivering a fundamentally more efficient and new approach to IT.

Zimbra is a leading open source e-mail and collaboration solution with over 55 million mailboxes. As an independent Yahoo! product division, in 2009, Zimbra achieved a mailbox growth of 86 per cent overall and 165 per cent among small and medium business customers.

BAE selects MontaVista over proprietary RTOS

The BAE Systems Bofors in Sweden has selected MontaVista Linux to power its new naval and artillery gun systems. The use of commercial quality embedded Linux provides for rapid development, combined with the long term support required for military applications. "Military customers demand systems of the highest quality, and require support for extended lengths of time. At the same time, BAE needed a development environment that allowed us to rapidly deliver new systems to market to meet the needs of our customers," said Mikael Alfredsson, manager electronics and electrical design, BAE Systems Bofors. "By selecting MontaVista Linux we are able to achieve all these objectives — an open source operating system of the highest quality, rapid development, and long term support from Linux experts."

Wipro joins the Open Handset Alliance

Wipro Technologies has joined the Open Handset Alliance, to extend its device development and service integration expertise around Android.

"Joining the Open Handset Alliance will help us turn around Android-based products and solutions faster with rich feature sets," said Nagamani Murthy, vice president, Mobile, Automotive and Consumer Electronics Group, Wipro Technologies. "With most device vendors now releasing Android-based phones, we foresee an impressive growth on this platform. The Open Handset Alliance membership will enable us to satisfy the strong customer demand," she added.

LFY Insight: Should You Bet on **Postgres**?

Rs 100
ISSN 0974-1054

LINUX

THE COMPLETE MAGAZINE ON OPEN SOURCE

ForYou

An EFYGROUP Publication
ISO 9001:2008 CERTIFIED

VOLUME: 07 ISSUE: 12 February 2010 116 PAGES ISSUE# 85

Experience Maemo 5

We take a look at [illegible] powered [illegible] point of view to understand [illegible] has matured over the years. [illegible]...

Other features inside

A Conceptual Introdu[illegible] to SCTP

Tried the WikiReade[illegible] It's Wikipedia in Your Pa[illegible]

HTC Tattoo Reviewed. It's Android for All

Backup Your Data with Back In Time

5 Tools to Develop AAA Games

Get Started with OpenLDAP

India	INR 100
US	$ 12
Singapore	S$ 9.5
Malaysia	MYR 19

The Wikipedia in Your Palm!

That's the WikiReader. A 127 gram electronic device with dimensions of 10x10x2 cm that contains the complete English language entries from Wikipedia.org.

When I was younger, the only places I could access the volumes of the Encyclopaedia Britannica were the school library and an uncle's place. Maybe I was just plain uninterested, or dumb—take your pick—I never fancied taking those hard covers out of the shelves to peek at what lay inside. Maybe all those fat volumes made me nervous, to begin with.

Fast forward to today: my dependence on Wikipedia is only next to Google. Whether I want to read up on musicians, places or software projects, Wikipedia pretty much has my requirements covered—that too, with enough citations to refer me to other places, if I'm interested.

So when the WikiReader [*thewikireader.com*] was announced a few months back—by the same folks who're behind the OpenMoko (Om) initiative—it was time for all Wikipedia fans (including myself) to rejoice. This device claimed to fit in more than three million English entries on the Wikipedia into a pocket-sized electronic device. What's more, it didn't require any form of connectivity at all—the articles being stored in a microSD card.

The great news is that we can now grab this device in India! Jaipur-based IDA Systems has started distributing it for Rs 5,500, inclusive of taxes and shipping charges [*www.idasystems.net/wikireader*]. If this seems like a steep charge, consider this: the Encyclopaedia Britannica costs Rs 63,000! Compared to that, the WikiReader seems quite easy on the pocket. Besides, a couple of night outs at a decent restaurant and you end up spending more than five grand. Still think the price is too steep? Anyway, since IDA sent a piece to us to review, I might as well get on with it.

The device comes in a no-nonsense black and white paper box with WikiReader written in bold—top, front side—and a picture of the device with the tagline "Wikipedia In Your Pocket" in the centre. Nice to see things being kept clean and simple.

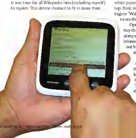

Open the box and take out the cardboard tray that holds the device inside a polypack, along with a 20-page manual with very basic information on how to get started. You can make out by the size of a manual how easy or difficult it will be to get started with a device.

As you can see in the picture, its user interface consists of a display screen that takes up most of the real estate, leaving just three rounded buttons at the bottom—search, history and random. Turn it on using another switch located at the top right-hand edge of the device. It'll take only a second to boot up and present you with the keypad that takes up around 40 per cent of the display's bottom portion for typing your search query—for example, 'Linux'. The top empty portion tells you to "Type a Word or Phrase". Easy, right?

This display is an 8.9-cm (3.5 inch) reflective monochrome LCD with a resolution of 240x208 pixels. And yes, it's a capacitive touchscreen, so you press the on-screen keypad with your bare fingers to enter the word/phrase you want to look for. Well this is where I came across my first obstacle. Either my fingers are too fat, or the screen space is too inadequate—I found myself often pressing the wrong characters. Also, you need to press a little bit harder to strike a note here. Makes me wonder, wouldn't a resistive screen have made better sense—I could have used a stylus.

Moving on, as you start entering the characters for your query, the top part of the screen shows you possible entries you might be looking for—it only shows five options, and there's no way you can scroll down to see if there are more matches. Could have been quite a handy option, don't you think? Upon entering Linux, it displays the following five options: 'LINUX', 'LinuX', 'Linux', 'Linux (disambiguation)' and 'Linux (GNULinux)'. All of them will take you to the Wikipedia entry on Linux, except for 'Linux (disambiguation)', of course.

So, I started to read the article on Linux, and I stumbled upon something weird right on the second word. It read: "Linux (commonly) is a generic term referring to Unix-like computer operating systems based on the Linux kernel ..." You'd wonder what's missing after the word 'commonly'. Well, it seems they have removed the phonetics. Or possibly they haven't, and something is wrong with the rendering that's stopping the characters from showing up. Since I can't really understand phonetics, it doesn't matter to me much. But still, it's something for the manufacturer to check on when a new update to the Wiki database is released.

The manufacturer also needs to come up with a way to represent tables somehow. In its current form, articles like ***http://en.wikipedia.org/wiki/Comparison_of_relational_database_management_systems*** only display the subheads, while the tables that are the actual content under the subheads are missing. It's understandable that inserting a table into this sort device with a smaller screen size is troublesome, and I've heard that Om is already working on a way to take care of this issue. So, a future update could solve this problem too.

Why updating all this is possible is because the whole content is stored in a microSD card. You can easily take it out by removing the batteries. So when the manufacturer releases an update to the database, you can download it onto your PC and copy it over to this microSD card. You also have an option to order updated cards by snail mail—they ship you a new microSD card for a fee. More details are here: ***http://thewikireader.com/update.html.***

Anyway, going again back to that Linux article, you might have noticed that I've left some words underlined while quoting the opening sentence. These are basically the same Wikipedia hyperlinks that take you to the Wiki article on the hyperlinked word/phrase. It works as expected on this device too.

What's inside the box

One thing I would like to point out is, when you click on one of these links, there's no back button to come back to the previous article you were reading. So, the only way is to press the 'History' button, which keeps a tab on all the articles you've previously accessed.

Talking about reading articles, the display resolution can hold a maximum of 13 lines of text with approximately 38 characters per line. The font size is adequate, making reading entries quite comfortable. What could be discomforting is the absence of any backlighting. So, you can only use this device at a comparatively brightly lit place. Also, since the screen doesn't support colours, it obviously doesn't display any pictures.

This pretty much sums up my brief report on the WikiReader. Did I miss something? Oh yes, it runs on battery power. Not some hard-to-find-in-a-general-store Li-Ion battery but two AAA batteries that are quite easy to find. Om claims the battery life is approximately 90 hours, which I'd say is quite decent. Besides, they have designed the Reader such that if it's inactive for a while, it automatically shuts off—a good power-saving technique. You might think restarting the device again is a headache. But considering it only takes a second to boot up, and a few more seconds to press 'history' and retrieve the article you were reading, it isn't much of a pain.

My verdict? Go for it! It's a handy present for your child or even for yourself, if you're a Wiki addict. Well, you have the 'random' button to take you to places—most of which I'm sure you wouldn't have visited because you never knew they existed.

By: Atanu Datta

He likes to head bang and play air guitar in his spare time. Oh, and he's also a part of the LFY Bureau.

HTC Tattoo

Brings Android to All

The HTC Tattoo, a bargain version of the company's flagship phones, is like an elderly HTC Hero—shrunken and less perky but still with plenty of fun times left in the tank. You'll definitely notice where they've cut corners—the resistive touchscreen is especially annoying—but, if you're looking to save some money, you can't beat this handset that offers all the power of an Android smartphone for just Rs 17,000.

The Tattoo offers the same powerful Android operating system and social networking goodies as the Hero, but in a cheaper, smaller phone. That means you can have all the power of a handset running Android—a usable, feature-packed smartphone OS—at a price that could normally only buy you a more stripped-down phone.

The Tattoo comes with a good suite of applications pre-installed, including Google Maps, Twitter and a YouTube app. If they're not enough, the Tattoo will also allow you to

▲ What's inside the box

▲ Home screen

shop in the Android Market, which hosts heaps of little programs that give the phone even more power. Fortunately, many of the apps in the Android Market are free, and it is easy to download and install everything from Spotify to Skype.

HTC has also iced the Android cake with its Sense user interface—the same one as on the HTC Hero that I raved about back in November. There are seven home screens, which you can customise to your heart's delight with wallpapers and widgets. These include Twitter and e-mail widgets that show your missives as they pop up in real time, without requiring you to open an app, like you would on the iPhone, for example.

You can also change your whole configuration depending on your mood, by creating your own themes, known as 'scenes'. For example, you may want to look at your work calendar during the week, but avoid it at all costs during the weekend. Your 'weekend' scene might show your music-player widget instead. But the phone is also well set up right out-of-the-box, so you don't have to change anything if you don't want to.

The look 'n' feel

While the Android OS is flashy and oozing glitz, the HTC Tattoo is rather unassuming.

The phone has a matt-silver finish and feels rather chunky. It's compact, though, and it looks sleek—with HTC disposing off the trackball as seen on the Magic and Hero, and replacing it with a flat, circular 'four-way navigational D-pad'.

In all, there are seven buttons on the chassis, broken into five sections, with the aforementioned navigational button being the most prominent.

You would expect, then, that this would be the main button to wake up the phone from sleep. Anyone familiar with the Magic will know that's not how it works, however, when it comes to operating the Tattoo. Instead, you click the much smaller menu button twice to get into the main interface. To close things up, simply press the power button on the bottom right.

If you're under-whelmed by the HTC Tattoo's demure form, then don't fret as the phone is all about individuality. HTC is releasing a number of funky cases with tattoo-like designs, so you can customise your handset. This is a good call on HTC's part, giving you the chance to choose your phone's style, rather than the company opting for a one-style-fits-all device.

When you flip the phone over (both the back and the front are customisable), what sticks out is the 3.2MP lens for the phone's camera. Unlike the Hero, which sported a small lens with a big 5 mega pixel count, the lens on the HTC Tattoo sticks out like a sore thumb, surrounded by a chunky metallic border.

The only other things to note on the chassis are a 3.5mm headphone jack, a micro USB connector, and volume control on the left-hand side, which is pretty much flush with the chassis.

When it comes to what you get with the phone, there are hardly few things in the box, except for a charger (whip the three-pin plug off and you get your standard USB data cable for the mobile as well) and a headphone

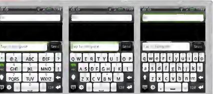

▲ Virtual keyboard options

Processor and memory

HTC Tattoo is powered by a single Qualcomm MSM 7225 processor running at a speed of up to 528 MHz. Overall, I havent experienced any lock-downs, crashes or lags. Moreover the phone responds very fast to any command, even when multiple applications are running in the background.

Tattoo has 256 MB RAM and a 512 MB ROM memory. The storage space can be expanded up to 16 GB through the hot-swappable microSD card slot that has been placed under the back cover, on the left, lower corner of the phone.

Calls and contacts

The core benefit of any phone is its ease in making calls. In Tattoo's case, it's very easy. On the home screen, just click the Phone tab and you are into your contacts database.

The dial pad is the first thing you will notice when on this screen. This isn't just there to type in full numbers; you can use it as a quick link to the number you want. Simply tap the number with the letter that your contact's name starts with, and it will jump to that section.

This is something that doesn't take too long to get used to. If you do get annoyed with it then just hit the grid button on the bottom left and you will have full scrolling control on all your contacts. Here, you can scroll through and choose a number. If you dont want to use the touchscreen, than the HTC's four-way navigational button can also be used to scroll down.

Hold down your finger on a contact and you will be able to see what messages have been sent to you by that person and other snippets of information.

The call quality is crisp and clear, and I don't have any complaints about the voice being distant or muffled. Turning up the sound on the flush set volume rocker on the phone's side did make a difference, but it also added a slight crackle. This didnt happen all the time and could well be blamed on the reception.

As long as you remember to press the *End Call* button on the screen, you shouldn't have a problem ending a phone call. It sounds like an obvious thing, but it's not, as you would instinctively want to press the angry red button on the right, which will put the phone on standby immediately.

Messaging

Text messaging is definitely an experience here. You have two options when texting people: you can either do it when the phone is in portrait or in landscape mode. I would recommend the latter. My fingertips arent the pudgiest, but when I used the phone in portrait, it took a long time to text someone, as I had to constantly correct what I had written.

Touchscreen phones are renowned for being a learning experience for even the most seasoned of texters. And in the case of the Tattoo, it is tougher since the screen is 1.016 cm (0.4 inches) smaller than the Magic and almost 2.5 cm (an inch) smaller than the Hero.

This makes a massive difference when it comes to texting, as even though there is an option to have the full QWERTY in portrait mode, the letters are far too crammed for you to use it properly—unless you have a sniper's aim.

Landscape mode means that letters have a bit more space to breath and fewer mistakes are made.

The Internet

As with Android phones, the HTC Tattoo is fully equipped to be an Internet browsing device. Getting on to the Web is fairly straightforward—just choose the Browser app and enter the URL.

Web access is fast and loads up pretty quickly even on the pathetic EDGE connection in Bangalore. And as far as the Wi-Fi network goes, even a paltry 256 KBps connection is good enough to give you a rich Web experience.

Once you are on the Web, the problems with resistive screens soon appear. There is no pinch and zoom on the HTC Tattoo and navigating your way around takes a little bit of added pressure.

A zoom option is at the bottom of the screen at all times, however and though there is more jitter than expected in this feature, it is definitely welcome. And there is also the tried and tested double tap feature to make the text on the screen fit the screen—nice and big to read.

When it comes to the hallowed Flash video, while the HTC Tattoo is described as capable of playing Flash, it didn't seem to do so, in my case. The YouTube videos, despite having been selected on the browser, got opened by the bundled YouTube player for streaming video.

When browsing a video, press *Menu* and there are a number of options to help your Web browsing along. Included is *Copy* and *Paste*, the ability to share a page and information regarding your viewing history.

By default, the phone takes you to mobile versions of the websites you wish to view. If this is something you don't want, then you can also turn this option off by digging deeper and checking out the Web settings on the phone.

Connectivity

The HTC Tattoo has plenty of ways to access the Internet, utilising both HSDPA and Quad-band GSM/GPRS/EDGE connectivity.

The phone also features seamless Wi-Fi access and always managed to log into an authorised Wi-Fi network every time the phone was within the range. The switch between EDGE and Wi-Fi was absolutely faultless and hardly showed any signs of unreliability even when I was watching videos on YouTube while making the network transition.

And to make things that little bit easier, you can also wade through the options and connect to Facebook, Twitter and Flickr.

I wish it were this easy to connect the phone to the PC. Unfortunately, there is no software inside the box. Instead, the phone asks you to download the software from HTC's website. Of course, the desktop application is absolutely useless for a hardcore Linux user.

An easier way to set the phone up was to go to *Data Synchronisation* and choose from *Google*, *Exchange ActiveSync* or *PC*—depending on which data you want to synch with the phone. This seemed to be an easier way to connect the phone to the computer.

A simple way to add music and movies to the phone was to choose the *Disk Drive* option once I was connected to the PC. This allowed me to freely drag and drop files into the phone.

The Tattoo may not quite hold its own against the likes of the Hero (its screen is to blame for that) and it will never take on the big boys when it comes to media playback, but for sheer functionality and ease of use, it's a winner.

Multimedia

Entertainment is the core requirement of mobile phones these days. Especially handsets that are targeted at the young generation need to have sound entertainment and multimedia highlights. Tattoo fulfils this requirement. There is a built-in music player in the phone that's easy to use and manage. Users can store music files according to their preference and sort them out as per the name of the song, artists, genres, etc. The music player is capable of playing MP3, eAAC+, WAV and WMA9 music files. The device also features stereo FM radio with RDS.

Sound quality is more than passable. Bass on the music I listened to was deep and sound clarity good. It's disappointing though that there's no equaliser. Moreover the earphones shipped with the phone were a terrible let down. So, in case you plan to buy the phone, I strongly recommend that you invest in a

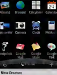

▲ Menu Structure

MobileBurn MobileBurn.com - Light version of Google Voice offers voicemail service for existing numbers: Google Voice new let...

▲ Twitter app

▲ Weather applet

better pair of headphones.

HTC Tattoo features a 3.15 MP auto-focus camera that takes snaps at a 2048 x 1536 pixel resolution. Video photography is also possible with this camera. Users can download or transfer video files and play them in the integrated video playback. This is capable of playing video files of different formats such as MP4, WMV9, H.263 and H.264.

A nice little feature is that when you lock your phone, you can still fool around with your music as much as you want, without the hassle of going into the Tattoo's OS.

Battery life

After some rigorous playing with the Tattoo—which included downloading many apps, playing myriad videos and songs, and a lot of Web surfing—the battery held up pretty decently. While the actual service time of the battery was nowhere near HTC's claim of 350 minutes, it wasn't too bad either.

One of the things I found was when charging the phone on the computer, the handset did get very warm near the micro USB connection. This happened on more than one occasion, but it soon cooled down once I stopped charging it.

Another excellent little add-on comes courtesy Android 1.6, and that is a battery indicator that gives you a break down of what your battery is being used for. This is a great way to keep track of whether or not you have closed down those unwanted apps.

And there are plenty of organising apps to sap your battery power if you so wish. There is a mixed bag of goodies on offer, including Google apps like Mail, Maps and Talk, and options including a calculator, calendar and FM radio.

I like it… very much

Ultimately, if the Android OS is ever going to win mass-market appeal, it is phones like the HTC Tattoo that will make this happen. Packing everything that is good about Android into a lower-cost handset just proves how good Google's OS and marketing skills are. And HTC has really made its mark re-skinning the OS to make it its own, and once again the HTC Tattoo proves that the Sense UI is one of the best things to happen to Android since its inception.

The compact form of the phone and the stylish design are among the best features of HTC Tattoo. Moreover, the wide range of high-end features makes the phone a must-have for Android fans who cannot buy themselves more expensive smartphones, such as the Hero or Magic.

I like it… not so much

The screen was lacklustre, with nowhere near the amount of zing that you get when using a capacitive rival. And when it came to the camera capabilities it was just, well, average. I think an accelerometer and multi-touch features would've been great to have. Also, perhaps Flash for the browser and more features for the music player.

It also disappointed when it came to text messaging. It's just not feasible to text quickly when the phone is in portrait mode. You will need to flip it into landscape and even then the predictive text will try and be evil—it feels like it works against you, not with you.

Another problem with the HTC Tattoo is that it seems to have adopted the same predictive text model as the HTC Hero. And just like Hero, it needs quite a few hours tinkering before you manage to discover where you can disable the feature.

On the whole the system is fine, but unless you are constantly looking at it, it will give you predictive words that you don't want. Keeping an eye on this and making sure that you are hitting the right letters is a bit of a pain.

It also has to be said that as the phone uses a resistive screen, you have to be even more precise in the hitting of letters and apply a bit more pressure than you would with a capacitive screen. I know that this type of screen has been used as a cost-cutting measure, but it does feel like a backward step for HTC, as it used capacitive screens on both the Magic and the Hero.

The verdict?

Frankly, I am in two minds about the HTC Tattoo. On the one hand, it's a likeable, compact Android smartphone running the cool Sense interface. On the other, the concessions the manufacturer has had to make to keep the price low are neither minor, nor just a few. I don't really have a problem with the screen size itself, but the lower resolution is equivalent to extremely poor catalogue of applications compatible with the device. This is a bitter pill to swallow and definitely something one is unwilling to put up with. More so now, when Android seems to be making evolutionary leaps forward and its development is picking up speed.

The unhandy resistive technology incorporated into the display and the barren waste that the application catalogue for the device is, paired with all the other minor yet real drawbacks of the handset, would have made the HTC Tattoo a viable option if it was available at an affordable price. Sadly, this is not the case. In case you are flat out determined to get a cell phone delivering the handy features of the socially-oriented Sense interface, I would definitely advise you to consider the HTC Hero—its price is not that much higher now and the handset is head and shoulders above the Tattoo.

By Sayantan Pal

More popular in the online world as Shayon, he runs a review blog on *Howzat.com* and also works as "the SMM guy" at Pronto Networks. Catch up with him on twitter (@shayonpal) or shoot him an e-mail (me@shayonpal.com) anytime you want.

Niyam Bhushan

Forget Linux

The future of our digital freedom is clouded

How *muft* and *mukt* will you find yourself in the next 10 years? Here's my bold prediction: almost all the software you use may become *muft* and *mukt*, but you will find your digital freedom almost vaporised in this new decade. You will become a digital slave with *muft* and *mukt* software.

Everything else

I attribute your imminent digital slavery to the overwhelming success of the Free Software Foundation (FSF) and to the FOSS movement at large. Their mandate has been to fight for the freedom and openness of software. Alas, software are not people and hence the irony. Some may debate that the adoption of Open Standards in computing is therefore even more important than just the adoption and development of FOSS. Standards aren't people, either. The ground-breaking initiative at CreativeCommons.org brings freedom to knowledge, education, entertainment, and all other forms of authored works. Not necessarily to people. Nevertheless, this decade will see an exponential surge in the use of CreativeCommons.org licences. Bringing digital freedom to people is another dimension altogether and one that's not yet been touched.

"Man is a social animal, with software to boot."

The nurturer vs defender

Some of you may point at the Electronic Frontier Foundation affectionately known as eff.org. According to its website, "EFF is the leading civil liberties group defending your rights in the digital world." Do consider spending at least 15 minutes browsing through the Web pages to see how this outfit fights Bluetooth and nail, taking on governments, multinational corporations and cartels who trample on your rights. Most of EFF's work and expertise is focused around the United States, though the EFF does strike at beach heads in several other countries from time to time.

Indians are in dire need of a formidable and indigenous EFF. If you're one of those helpless Indians frustrated in the face of spam marketing calls and text messages, or have been victimised by inflated electricity bills due to a computer error, you'll know why. In fact, it's about time some strong and committed individuals got together to launch EFF or something similar in India, within this year.

By its own admission, EFF is a defender of your digital freedom. Which means that it may choose to take up your case legally, file public interest litigations, and run advocacy programs. What I dream about is not an organisation and a movement that fights for your digital freedom, but nurtures it. Am not looking for the doctor in the emergency ward, but the yoga, fitness, ayurveda and meditation guru, all rolled into one—to help me live a healthy, fulfilled, peaceful and blissful life.

Man is a software animal

This think tank needs to address questions and find solutions to problems in this next wave of computing. In the previous wave, man was merely a software animal. In this wave, man is a social animal, with software to boot. So, if my data is stuck in the Internet cloud, the value is not in my tweets, blogs and comments alone, but the sharing, re-tweeting, commenting, and collaboration that happens around it. I could get the entire Facebook engine as copyleft software, run it on my own FOSS-based servers, but how do I get the social wealth embedded within Facebook, a wealth of which I too am a part, to which I've also contributed? Who owns this? Facebook, or all of us? Do I own my recommendations and reviews on Amazon.com? If so, do I get a share of the revenue they help generate for Amazon.com? These questions are one of many to be raised with Google services, Twitter, Flickr, Last.fm and everywhere else you dwell on the Web as a social animal. I haven't even raised questions focused around mobile phone software and services, where all consumers are being taken on a solid telecommute-ride by service-providers and their cronies.

All these questions are deeply thought provoking. These next 10 years are going to be the most exciting, challenging and stimulating times in the history of computing yet. Am buzzing and bursting with ideas, and look forward to sharing your ideas and suggestions as we explore unchartered territory, month after month, in this new avatar of Freedom Yug. Here Be Dragons and Rakshasas. Bring me my bow of burning gold. Bring me my arrows of desire. Bring me my spear. O Web clouds unfurl. Bring me my chariots of fire. END

About Niyam

Inspired by the vision of Osho. Copyright February 2010: Niyam Bhushan. freedomyugs at gmail dotcom. First published in *LINUX For You* magazine. Verbatim copying, publishing and distribution of this article is encouraged in any language and medium, so long as this copyright notice is preserved. In Hindi 'muft' means 'free-of-cost', and 'mukt' means 'with freedom'.

STOP

SANIsoft Technologies

An Open Source Map that Leads to the Treasure!

Show me the money? Is the question thrown at every open source entrepreneur who encourages fellow entrepreneurs to follow suit. Nagpur based SANIsoft Technologies reveals some smart open source business tactics.

From college classrooms to office cafeterias and family dinner tables—Phunsukh Wangdu has been the topic of discussion ever since *3 Idiots* hit the box office. His saga of making it big from modest beginnings, coupled with his strong passion, boldness and constant strive towards excellence left everyone in awe of him. We often bump into variants of Phunsukh Wangdu but seldom realise it.

Entrepreneurs in the open source business, brimming with passion for free and open software, are Phunsukhs in the making. Their courage to open the source code of their products and solutions speaks for their confidence and focus on excellence. Yet, doubts remain on the feasibility of translating their passion into a commercially successful business. Tarique Sani, CTO of Nagpur based SANIsoft Technologies, along with his wife, Swati Sani, erased all doubts on that front by building a viable business model around their Web development company doing PHP programming with LAMP. Starting out small from Sani's house in 1998, SANIsoft grew 100 per cent, year-on year for the next five years, multiplying its revenue by 32 times.

"We began very small, with just our brains and one computer. Back then, the Internet as well as Web development was new to India. We didn't really do market research to see if we would fit in but just believed in our programming skills and sold our expertise," says Swati Sani, executive director, SANIsoft Technologies. Today SANIsoft is growing at a rate of 40 per cent, year-on-year.

Building blocks

To create a niche for itself SANIsoft's initial business strategy focused on selling solutions based on open source platforms rather than open source solutions, and that became its key differentiator. Further, the team focused on making the most of its team members' skills. While Tarique Sani, with 16 years of programming experience, focused on getting the solutions right, Swati Sani, who had been a communications specialist for nine years, managed the marketing. SANIsoft also contributed to open source initiatives. "We actively participated in the open source e-groups, wherein the open source code we wrote spoke for the quality of our work. We also wrote technical and semi-technical articles for some websites and magazines," says Sani.

This three-pronged initial business strategy finally paid off. "Our first customer came to know of us through one of the e-groups we managed. We developed a classified and coupons application for Avinash Mohanl (currently CEO and president of NetCliff)," says Sani. Sugartrade, KlubQ, Sunflag Iron and Steel, and Puma Herbals were some initial clients.

However, making the business click wasn't hassle-free for SANIsoft. Being an adolescent firm, prospective customers had their reservations. "The going was tough considering the newness of the platform we chose to operate on. Initially, we were able to close less than 1 per cent of the deals. But persistence, active participation in the community, self-belief and our chosen technology soon paid dividends," says Sani.

Finding the appropriate talent to channel growth was another challenge. "Hiring was very difficult as most candidates weren't familiar with PHP. We picked up freshers and had to invest in training them in-house. All our current senior managers joined us as trainees nearly six to seven years ago and were trained in-house by our CTO," says Sani.

Consolidating initial successes

After sidestepping initial hurdles, SANIsoft soon created a niche in high performance PHP applications, moving from catering to small start-ups to bigger enterprises. Today SANIsoft has applications like Cheesecake Photoblog Software (a feature-rich photoblog software written using CakePHP framework), Coppermine Picture Gallery (a popular self-hosted photo gallery with six million downloads from Sourceforge), WAPpop (which gives POP access to glance through an e-mail even when on the move), PHPlib (a library of PHP classes which added sessions to PHP version 3), etc, in its portfolio. The firm's clients include Asiarooms (an Asia Pacific based travel house), Aximize (a property management portal), Budget Pulse (an interactive Web based budgeting utility), Sandiago Properties (a real estate classifieds agency), etc.

But how did SANIsoft create this niche? It was the cumulative effect of its business model and continuous community work that worked for the company.

Strong business model

SANIsoft has adopted a hybrid business model. Most of its revenue is generated by the 'loss leader' model, wherein a no-charge open source product is used first to attract the customer. This generates little or no revenue, but the product woos customers into buying other products, solutions or services of the company that are sold using the traditional software business model. "Loss leader means a software product is initially open sourced. The company derives a marketing benefit by way of showcasing its capability and code quality through this open sourced product and then procures business. SANIsoft primarily uses this business model. For instance, our open sourced Cheesecake-Photoblog based on CakePHP framework was cited by many as a reference for a CakePHP application. It helped us showcase our work, attracting business opportunities," says Sani.

The team also used the 'sell it, free it' model, wherein the software starts out as a traditional commercial software product but is converted to an open source product at an appropriate time. "A lot of development that we did for a company that partnered with Flickr for photo printing found its way into Coppermine code. Zip Locator, too, which was initially developed for classifieds was later freed. Similarly, B2B classifieds and B2C began life as in-house closed source projects. These have been our bestsellers, but since it was realised that none of the clients ever asked for an exact replica, we decided to open the source code of these products under the LGPL," says Sani.

A small proportion of the firm's business also comes from 'support selling', which involves revenue generated from media distribution, branding, training, consulting, custom development, and post-sales support. "We provide support for PHP-based open sourced solutions like Joomla, Coppermine picture gallery and Wordpress, apart from long-term maintenance and support for our own custom-made solutions for our clients," says Sani.

Active community work

Further, SANIsoft employees have been continuously involved in the community, focusing on bug fixing, submitting patches and replying to answers on mailing lists. This eventually gets the firm into the limelight, increasing its reach. "Giving back to the community is a deeply ingrained habit in SANIsoft. Our programmers help people with various problems in their free time or they develop free software. We have regularly contributed in covering costs during conferences for the KDE community in India. We have sponsored an international CakePHP conference in Argentina, and another in Germany," says Sani.

Looking ahead

Expansion is on the cards for SANIsoft. "We are in the process of doubling our programmer strength and thus increasing our turnover. We also plan to start offering turnkey and end-to-end solutions," says Sani. So, it seems like SANIsoft just might achieve a 100 per cent growth rate over the next five years.

By: Vanisha Joseph

The author loves to experiment and writing for LINUX For You is her latest experiment. So, beware! Just a minute, she also happens to be a journalist during the day.

An OS for Personal Computing

Let's take a look at the BeOS-inspired Haiku OS.

GNU/Linux has lots of features for the desktop and the server side. However, there are problems with Linux-based operating systems. Being a monolithic kernel, people often find the system becomes unresponsive when using a GNU/Linux system. Another major problem, especially for new users, is choosing between various distributions. They end up installing a distribution that has more apps and services than they actually need (or that their hardware can support) for their day-to-day use, which also serves to slow down their systems.

This article introduces you to an operating system called Haiku, which serves as a good starting point for aspiring students and those interested in hacking on operating systems.

So what's so special about Haiku? It is a descendant of the popular BeOS. The team behind BeOS had lots of ideas about making a perfect operating system targeted at personal computers. Unfortunately, it failed miserably due to various reasons.

Haiku is a POSIX-compliant operating system project that started in 2001. It was initially named OpenBeOS since it was heavily inspired by BeOS. Haiku specifically targets the personal computing space, to bring users a clean user interface and a highly responsive kernel.

It took almost nine years to get the first alpha release of the OS out on September 14, 2009. Some good features of Haiku are listed below:

- A specific focus on personal computing
- A custom kernel designed for responsiveness
- A fully threaded design for great efficiency with a multi-processor/core CPU
- A rich OO (object oriented) API for faster development

- A database-like filesystem (OpenBFS) with support for indexed metadata
- A unified, cohesive interface

Personal computing

BeOS was started with the aim to become a compelling replacement for Apple Macintosh systems. A great design and attractive user interface were part of the package. Haiku follows the same principles like the BeOS to deliver an open source OS with a simple, yet attractive UI, and supports x86-based hardware only. Already many applications developed for BeOS are available for Haiku and active development is on for new apps too.

Kernel

As indicated earlier, Haiku does not use the Linux or any other existing kernel. The Haiku project has built its own custom kernel using C++ (as the dominant) language. The modular hybrid nature of the kernel makes it highly responsive while running user applications on top of it. The kernel design also exploits features of modern day processors.

The filesystem

Haiku implements the OpenBFS filesystem that has superior design goals (and is POSIX compliant, of course). This is a 64-bit journaling filesystem capable of handling volumes close to 2 exabytes (1 billion gigabytes) approximately, with a maximum single file size of 260 GB. It is a database like filesystem that supports indexed metadata.

Install experience

The ISO image of Haiku is 380 MB and it's easy to install it on a hard disk using a few simple steps. The ISO can also be used to evaluate the OS in live CD mode. Here are some of my observations while playing around with Haiku:

1. A simple user interface (similar to Fluxbox or other light weight desktop environments)
2. There's a need to use the right mouse button a lot to navigate through the applications
3. It takes less than 10 minutes for a complete installation of the OS to the hard disk.
4. Currently it takes less time to boot
5. The required shortcut icons are present on the desktop.
6. It ships with the Firefox Web browser (BonEcho)
7. There are other applications such as a PDF viewer, paint utilities, document editor, etc
8. The tools required for the development of Haiku are included
9. Various applications and systems utilities are grouped in the menu as shown in Figure 2

However, being the alpha release, let's keep our expectations low. Some of the glitches in Haiku as of now which will be corrected in future releases, are as follows.

1. Installation does not provide a boot loader and partition editor
2. Graphics card support is not mature enough

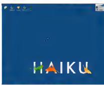

Figure 1: Haiku OS after installing it to hard disk

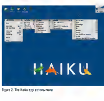

Figure 2: The Haiku applications menu

3. Wireless support is not available
4. Many hardware drivers are missing.

Overall, Haiku looks like a promising and simple OS. Those who are interested in C++ systems programming, and would like to contribute to a free software project, could get their hands dirty with the OS code base. There are more than 1,000 bugs open under various categories and already 3,000 bugs have been fixed (see reference [2] for details). The project also has a well established collaborative environment for development, ranging from mailing lists to bug trackers and version control. END

References

- http://www.haiku-os.org/
- http://dev.haiku-os.org/

By: Senthil Kumaran S

The author is currently employed by CollabNet, working for its Version Control Group. He is a full committer of the Subversion project and is a free software enthusiast. To know more, visit www.stylesen.org

Chrome Extensions Development

It's Fun and Addictive

The Google Chrome browser burst upon the scene a little over a year ago. Chromium is the open source project that Google Chrome is based on. Beta versions of Google Chrome for Linux and Mac were released in December 2009, along with extensions.

Extensions are small additions of features for the browser to improve its usability. Adding these is simpler than modifying the browser's source code. The vibrant Firefox extension community is one aspect that sets it apart from Chrome, so it will be interesting to see how Chrome's extension community develops.

Mouse Stroke is an extension for Google Chrome that lets you use mouse gestures to do things quickly within a browser—such as opening and closing tabs or moving back and forth between pages.

Here is a Q and A session with Quan Gan, author of the Mouse Stroke, where he shares his experiences with Chrome.

Q Tell us about yourself.

I'm a software engineer living in Beijing, China. I'm 28 years old. Although I'm not as involved with open source development as many others around, I do have some projects—mostly tiny ones—hosted at both SourceForge and Google Code.

Q How did you get interested in Google Chrome and its extensions?

I started to use Chromium because Firefox is slow on my Debian box. I found Chromium breezing fast, and started using it under Windows too. I started to make extensions for Chromium because I wanted to port my GreaseMonkey script, a telnet emulator, to Chromium so that I could secretly browse a telnet BBS at work :)

I didn't really get interested in developing mouse gesture extensions at the beginning. I was happy with StrokeIt under Windows and Easystroke under X. However, my laptop at the office is strictly monitored by my administrator—I'm not allowed to install any software—so I decided to develop a mouse gesture extension.

It's called Mouse Stroke. I posted the first version at *chromeextensions.org* and said to myself, "That's it... the final version!" But it turned out that people hated it :-)

A bug in the first version killed many websites, including GMail. So I had to pick up the source and make another version, then another version. At least it doesn't kill GMail anymore :)

Ironically, I started to use Chromium because it's fast on Linux, but my extension still doesn't work on the Linux version of Chromium. That's because Linux Chromium handles mouse events differently from the Windows version, which makes it impossible for the developers to control the context menu. This also applies to the Mac

Mouse Stroke in action

Mouse Stroke options

version. There's a reason why there are only a few gesture extensions that work on Linux and Mac—or are there enough? I'm hoping this can be solved soon but, well, I don't know when.

I'm currently working on user-defined strokes. Most gesture extensions on Firefox have this feature. The idea is to let users create their own gesture handling scripts. Hopefully it will be available on Mouse Stroke soon.

[Update: Mouse Stroke now works on Linux and Mac, and the feature for user defined strokes is also available.]

Q Do you work on software for a living? Would you call yourself a developer or a Web designer?

I'm a developer. My work is to develop software to process financial data. I was interested in Web development back in college, mainly using PHP. I wrote my own forum software for my WoW guild, a website for my lab. I also developed a website for someone in Boston years ago. But I don't think I'm a Web developer. Most of my programs are written in C++, PL/SQL, and recently Perl. However, Web development is fun—I'm spending more and more time on it.

Q Chrome extensions are only JS, HTML, and CSS + images, right? And it sounds like you are more of a traditional application developer. So is it easy to work on an extension instead of using C++ or Perl?

Chrome extensions are JS, HTML and CSS, but it's not regular JS that's used on websites. Google implemented some APIs to expose the internals of Chrome, so extensions can do things like add bookmarks, move tabs around, etc. However, this API set is still minimal—there's still a lot of work to be done to make Chrome's extension system as powerful as Firefox's.

Developing browser extensions (or Web applications) is easier than developing traditional applications, because recent browsers are all packed with great developer tools. All you need is a browser and a text editor. But for traditional languages, you have to know the compiler, the linker and the debugger. You have to know different runtime environments. There are so many things to take care of.

Besides, Web/extension development is more fun and more addictive. We rely on Web applications so much these days and a simple browser extension could help you immensely. Also, users can get access to your applications in a much easier way. END

References

- Google Chrome: *www.google.com/chrome*
- Chromium project: *www.chromium.org*
- Google Chrome announcement: *chrome.blogspot.com/2009/12/google-chrome-for-holidays-mac-linux.html*
- Google Chrome Extensions: *chrome.google.com/extensions*
- Mouse Stroke Extension: *code.google.com/p/chrome-mouse-stroke*

By: Anthony Bryan

The author loves the arts and technology, and is trying to improve downloads with the Metalink project.

OPEN SOURCE INDIA
Chennai Trade Centre
19 21 September 2010
CHENNAI

WE SEEK YOUR SUPPORT

to make an Indian Open Source event a global phenomenon

We invite speakers, volunteers, subject-matter experts, LUGs, media partners, exhibitors and sponsors to engage with us and help Open Source India become an event that makes all other nations take note

KEY FACTS

Dates: September 19-21, 2010
Venue: Chennai Trade Centre (Chennai, India)
Edition: 7th (commenced in 2003)

Here's How You Can HELP...

Share your expertise by speaking on any of these topics...

- Developing Web Applications Using Open Source
- Developing Mobile Applications Using Open Source
- Developing Embedded Systems (or Device-Level Software) using Open Source
- Managing IT Infrastructure Using Open Source
- Migrating IT Infrastructure to Open Source
- Business ready Open Source Applications
- Why Open Source spells 'moolah' for ISVs
- Cutting IT Costs Through 'Open' Cloud Computing & Virtualization

IT organisations can...

- Offer senior personnel as member of our 'content planning & review' team
- Join as 'knowledge partners' by sponsoring speakers
- Join us as 'delegate partners' by sponsoring bulk delegates

Anyone can...

- Be there as a delegate
- Be there as a volunteer
- Put OSI logo on your blog, Facebook, Twitter etc
- Register on Open Source India's website for updates
- Become a fan of Open Source India's Facebook page
- Follow @osidays (on Twitter)
- Spread the word about Open Source India
- Send us suggestions and feedback

LUGs can partner with us, and...

- Get special benefits for their members
- Get table space to demo open source projects
- Highlight efforts made by them amongst the audience
- Organise independent BOF sessions and workshops

Commercial organisations can...

- Exhibit within the expo zone
- Sponsor our sessions, tracks or the entire event
- Sponsor travel and stay of 'star speakers'

Become our media partners

- Write about Open Source and spread awareness
- Get journalists to register for 'Press Passes'

VITAL STATISTICS

Goal:
To promote open source in India.
Started in:
2003 (branded as LinuxAsia)
Last Edition:
Open Source India 2009,
held at Chennai in Feb 2009

Break-up of attendees (2009):
Total techies: 2500, CXOs: 150,
Senior Academicians: 150+, IT Managers: 250+,
Software Development Professionals: 700+
Status
#1 open source event in South Asia

You Can Contact Us At:
WEB: www.osidays.com | **EMAIL:** info@osidays.com | **TWITTER:** www.twitter.com/osidays **FACEBOOK:** Open Source India' fan page
PHONE: 011-26810601 or 02 or 03 or 09811206582 | **ADDRESS:** D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110020

Powered By

Media Partner

Who Needs Just a Smartphone?

Experience Maemo 5

It's a phone...It's an Internet tablet...
It's a portable computer... It's a shark?
Well, just call it the Maemo powered N900!

Well, to begin with, I'm filing this review after playing with the phone for just three days—the deadline's up on me and Nokia India didn't have an N900 at their disposal to send us a device any earlier. So, you ask me about my impressions? Hmmm...

Let the truth be out that I'm yet to try an Android phone. Neither have I tried any of Nokia's earlier-generation Maemo-powered tablets. So, am I qualified to review this device? I'd say it really depends on what you expect from this review and more specifically, what you expect N900 to be.

The buzzwords these days are smartphones. Or is it the recently-announced "superphone"? I'd say who needs a phone anyway? I'd rather be happy if it's a portable computer first and a phone second, rather than the opposite. And if we think alike, you should take a hard look at the Maemo-powered N900, before you go all ga-ga over the Google Nexus One because of its awesome hardware specs. Because N900 is a different breed all together.

The device

The N900 is a rectangular brick with the dimensions 110.9×59.8×18 (19.55 at the thickest part) mm that weighs approx 180 grams. The front of the device only has the screen—an 89 mm (3.5-inch) touch sensitive resistive widescreen display with 800×480 pixel resolution. And the landscape should be your default view rather than portrait that we've come to expect from all phone-like devices. Well, the N900 is a tablet first, anyway.

Holding it in a landscape position you have a few things on the extreme left side of the N900's front. There's a light sensor, a VGA camera (for video conferencing), proximity sensor, earpiece and the indicator light. But no physical buttons (not even one to take calls) that we expect in phones.

On the top you have the volume slider, the power button and the camera button. Flip to its right, and you get one of the stereo speakers, the lock/unload key and the 3.5-mm headset jack. You also have the stylus nicely tucked inside at the bottom—a very convenient position I'd say since you'd generally want to take it out using your right hand. On the left, you have the second speaker, the mini USB port, that doubles up as a charging point too. I would have preferred a standard mini USB port, so that we could make use of regular mini USB cables to connect the device to a computer.

The three-row keyboard slides out from the bottom. The keys are slightly raised so that you can comfortably use the keyboard without looking at keys once you're used to the location of all the alphabets and special characters. These special characters are placed at weird locations I must say but that's quite natural for a portable device. And I'd say Nokia has done a good job in designing this keyboard layout.

▼ Carl Zeiss 5 MP camera lens

The back side of the phone has the 5MP camera, which is protected by a lens cover. Open the cover and you'll find a dual LED flash next to the lens. Around the lens cover is the N900's stand.

The device is powered by an OMAP 3430 (ARM Cortex A8) 600 MHz processor, PowerVR SGX with OpenGL ES 2.0 support for graphics and 256MB of physical RAM (plus 768 MB virtual memory). It comes with a 1350mAh battery and 32 GB storage (which you can expand by up to 16 GB by hot swapping a microSD card).

Other accessories include a headset (which I found not really comfortable to wear but has good sound quality and output), a USB cable, a charger and a video out cable. And of course, there's a manual and a cleaning cloth, too. Nice! Well, here's the complete spec sheet, *http://maemo.nokia.com/n900/specifications/*

Overall, the device looks nice and rugged, and a bit on the heavier side—not at all sleek like the iPhone or the Nexus One. But, heck, it's more like a power device than just a smartphone. And the looks, I'd add, are a must for the geek who uses it. Besides, there's a complete GNU/Linux distro that powers the device, and not just a Linux kernel and a few libraries to support a JVM.

Getting started with a flash

Okay, let's put it this way. I didn't like the fact that I was stuck with an older version of Maemo [Version 2.2009-47-23.002] when an update was released on January 14 [Version: 2.2009.51-1, details at: ***http://wiki.maemo.org/Maemo_5/PR1.1***]. It's supposed to work over the air as in the package repository should have notified me—which didn't happen. Upon refreshing the repository multiple times it would still inform of no updates. So, I headed straight to download the firmware and flashed the N900. Details on how to do it are here: ***http://wiki.maemo.org/Updating_the_firmware***.

The best part about flashing the device was, it was a piece of cake using a Linux desktop also. I'm so glad that, for once, something didn't force me to use a PC Suite that only works on that OS-that-shall-not-be-named. The whole process took hardly a few minutes, and I was back on track.

Note: The reason the update didn't happen over the air could be attributed to a fact that I'll explain later

NVIDIA® Tesla™ BIO WORKBENCH ENABLING NEW SCIENCE

Tyrone®

The NVIDIA® Tesla™ Bio Workbench enables bio-physicists and computational chemists to push the boundaries of bio-chemical research. It turns a standard PC into a 'computational laboratory' capable of running complex bio-science codes – in fields such as drug discovery and DNA sequencing – more than 10-20 times faster through the use of NVIDIA Tesla computing processors. Molecular simulations that had been possible only by using supercomputers can now be run on an individual workstation, optimizing the scientific workflow and accelerating the pace of research.

Life Sciences Applications Available on CUDA

- Amber 10
- NAMD
- VMD
- GROMACS
- LAMMPS
- GROMOS
- TeraChem
- aceMD
- HMMER
- GPU-AutoDock
- MUMmerGPU
- HOOMD blue

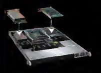

- 2 Teraflops in 1U (240 cores per GPU)
- The Most highly optimized multi-GPU systems available
- Gold Level 93% efficiency 1400w power supplies
- Flexible system configurations with up to 4 GPUs

- 4 TFLOPS in Tower
- Up to 960+Cores
- PSC: Personal Super Computer

To learn more about bioscience applications accelerated on Tesla GPUs
http://www.nvidia.com/object/tesla_bio_workbench.html

Powered by SUPERMICRO

Do call or write to us for more details

Netweb TECHNOLOGIES
An ISO 9001:2008 Company

Corporate Office
B-1/A25, 2nd Floor
Mohan Co-Operative Industrial Estate
Mathura Road, New Delhi - 110044

Website : www.netwebindia.com
www.tyronesystems.com

Email : info@netwebindia.com
netweb@vsnl.com

New Delhi	+91 11 29942258/29942277
Bangalore	+91 80 41126966/65560930
Chennai	+91 44 43531009/26631975
Kolkata	+91 33 22232085/23232956
Mumbai	+91 22 40159811/40158815
Singapore	+65 64157239/92289638

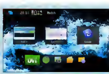

Figure 1: Workspace #1

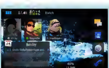

Figure 2: Workspace #2

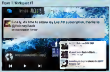

Finally, it's time to renew my Last.fm subscription, thanks to @BaloneyGeek

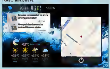

Figure 3: Workspace #3

Figure 4: Workspace #4

The Maemo desktop

Back on the Maemo (shall I call it) desktop, it asked me to enter my regional settings. I had to enter English UK there... Later from settings, I could change the time to IST but not the region to India :-(

On the desktop, things are in a landscape mode—the way a desktop should be—and not portrait that you expect in a phone. The resolution is 800x480. This means you have enough space to fit in a lot of things, and I mean widgets, shortcuts, et al.

Take a look at Figure 1. You can see three website shortcuts for Twitter, Facebook and Google. At the bottom, there's a tile for things provided by Nokia's OVI. Taping the first icon that says OVI opens the Web browser and takes you to *www.ovi.com*. This is where you can see all the services that OVI provides. The next icon, which shows a compass, opens the OVI maps application—if you're connected to the Internet, using GPS it can show your location using OVI Maps. The next three icons are OVI's online services: OVI share, OVI Files, and OVI e-mail.

You need to create an OVI account to access these features. Although I did create an account, I didn't find myself intrigued enough to try out the services. I'd recommend you to see for yourself whether you'd need these services or not.

Back on the desktop again, there are three more workspaces besides this one—yes, just like you have four virtual desktops on a GNU/Linux desktop. Simply swap your finger on the screen horizontally to go to the next one. Ain't it great?

So like on your desktop computer, customise these for your different sets of task loads. Figures 2, 3 and 4 show the way I've customised each of my workspaces:

- Figure 2 – Three favourite contacts (the ones with the pictures), icons for *Conversations* (IM, SMS, Twitter), *Contacts* (address book), *E-mail* and *Phone*; besides a widget that shows me the current processor and RAM usage.
- Figure 3 – Twitter and Facebook widgets, and another one to control my media player.
- Figure 4 – an AP (Associated Press) ticker to keep myself updated of the breaking news, weather applet and GPS location.

Customising the desktops is easy and very intuitive. Just tap an empty place on the screen and you get a drop-down from the top right corner with an icon of a gear. Tap that, and the top portion of the desktop changes into a horizontal panel with the *Desktop menu* (Figure 5).

Tap the menu, and you get a drop-down with different items—add an application shortcut, add bookmark, change the desktop background, change the desktop theme, add a contact, add widgets and manage views (which basically lets you select the number of views, and four is the maximum you can have). See Figure 6.

As you can also see in the screenshot, a very cool visual effect

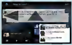
Figure 5: Customising the workspace

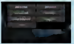
Figure 6: Options for customisation

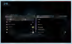
Figure 7: Dashboard shows two apps running

Maemo gives you is that when a menu pops out, the rest of the desktop screen becomes blurred. You simply tap on the blurred portion of the screen (that is, anywhere outside the menu) to get back to the desktop. Once the icons/widgets are on screen you can move them around with your finger and place them at the right positions.

As you can see in Figure 5, all these widgets/shortcuts also have an *X* icon on their top-right edge to close/delete them, and an additional config icon on their bottom left edge (if the item needs any further configuration). For example, you'll need to enter your authentication details into the Twitter/Facebook widgets to have it display the updates. When you're satisfied with their placements, tap the *Done* button on the top right portion of the screen.

The most interesting part of the desktop is perhaps the dashboard. Take a look at Figure 1 and 2 again. The blue icon on the extreme top-left corner of the screen is to access the dashboard. However the icons are different in both the figures. Why?

In Figure 1, where it shows two overlapping blue rectangles (windows), it indicates that there are applications running in the background. Tap on it to access these. Figure 7 shows how things are arranged when only two apps are open, while Figure 8 shows 13 apps open—well, multitasking is a piece of cake on Maemo as you can see.

The dashboard can fit up to 12 apps. If you open more apps, you'll have to scroll using your finger.

In Figures 2, you can see the dashboard icon replaced by an icon with a set of six small blue rectangles. Tap it to access the application menu (Figure 9). [To get back to the desktop, tap an empty section of the screen.] In Figure 2, I wasn't running any application in the background, and hence it showed the app menu instead of the dashboard icon.

As you can see in the application menu in Figure 9, the list consists of typical things you get in any phone. However, More... (the last icon on the list) is where the rest of the apps go. Tap and view the apps—Figure 10.

By default, there are about a dozen apps here—the screenshot shows a lot more because I'd already installed a lot of apps before I sat down to write about them. Out of the box, N900 comes with a few basic apps like a note-taking application, a PDF reader, RSS reader, a sketching app (so you can use N900 as a digital pen), a few games, an X terminal (yay!), file manager and application manager. Using this App Manager you can install more apps, configure repositories, etc. Yes, it's the *apt-get* magic. Did I tell you, Maemo is based on Debian? ;-)

Before I get into the details, let me again get back to the desktop where we have the system tray that sits next to the dashboard icon. It shows you the clock, the cellphone network strength, battery indicator, the Wi-Fi radar if you're connected to an access point, the IM status, GPS, etc—basically things that go into a typical desktop system tray. Tap it to access detailed statuses or configure things (Figure 11).

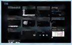
Figure 8: Dashboard shows 13 apps running

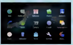
Figure 9: Applications menu

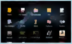
Figure 10: 'More...' apps menu

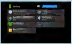
Figure 11: System tray details

Note that, I'm giving you a demonstration using a fully setup device. You'll need to set up a few things (even install a few apps from third party repos) before getting started, obviously.

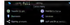

Figure 12: Settings menu

Setting it up

The first thing to do is check out what all you can configure—that's right, tap *Settings* (under the app menu). See Figure 12. Different things you can configure in the device are placed here under relevant categories like:

- Personalisation – It lets you set the themes, phone profiles (like ring tones, silent mode, vibrations, etc), region, date display, notification light, text input, etc.
- Connectivity – This has things like Internet connections (lets you choose and set preferences between WLAN, GPRS, 3G, etc), online sharing accounts (OVI, Facebook and Flickr), Bluetooth, GPS settings, phone settings, VOIP/IM accounts (besides OVI you have Skype, AIM, Google Talk, Yahoo!, MSN, etc), Bluetooth, etc. Heck, you even have an FM transmitter here :-)
- General – This one lists memory (storage space consumption and availability), TV out, screen calibration, device lock settings and *About product*. [With the help of this last option I was able to figure out that the N900 didn't have the latest Maemo 5 update.]

Things I took care of here at one go were: setting up the WLAN connection, IM accounts and text input. This last option lets you enable a virtual keyboard also (disabled by default), thus making you 'sort of' independent of the physical QWERTY keyboard. Although I love the physical keyboard, my wife seemed to prefer the virtual one—God knows why she liked using it with the stylus, or maybe it was just for fun.

Full-blown Web browsing

Well, once the WLAN's up, what else can an Internet addict do besides opening the Web browser? This one comes with a Mozilla-powered browser (called MicroB) and it's damn exciting to use—thanks to the default landscape layout and the keyboard.

Tap Web from the app menu, and you open the browser. At the bottom is the address bar and the top portion has your bookmarks with a thumbnail view of the page (Figure 13). Enter an address and hit the return key—it opens the website in a new window leaving this home window in the background.

Seriously, it feels just like I'm using a regular full-blown browser—no need to go back to things like Opera Mini. Most websites open in their default layout (not mobile) and because of the larger screen resolution, and the landscape mode, you needn't horizontally scroll to read, say, Wikipedia entries (Figure 14). Once a website loads up, the browser automatically switches to the full-screen mode to give you that full 800x480 px resolution for your reading pleasure. You get a small icon on the right-bottom portion of the screen to go back to the windowed mode.

For a quick and casual zoom in, simply double tab on the Web page—another double tab reverts to the original size. If you need finer control over zoom, use the hardware volume controller. And thanks to the availability of Adobe Flash plug-in, YouTube and other Flash videos play the way they're supposed to, inside the browser and not on some separate app. See Figure 15.

A brilliant feature is the ability to bring in the mouse cursor through a simple gesture. Swipe your finger fast (or the stylus) from outside the left-side of the screen towards right (that is, into the screen), and you get a small mouse icon. Tap it and you get the the cursor. Now simply drag the cursor on text to select it (Figure 16) and hit Ctrl+c to copy—yes, you can paste it anywhere using Ctrl+p.

Another nice feature is when you press your finger gently on the screen, a pop up menu appears with the options to open a new window, add the Web page URL (as a bookmark or a desktop shortcut) or "save page as" to save it somewhere in your filesystem.

When in the windowed mode, tapping the title bar will give you another menu. Some of the options are a repeat from the tap gesture menu, but you also have things like reload the page, search on page, FF download manager, Web page details, and browser preferences (that is, *Options*). From *Options* you can configure the browser settings, or access Add-ons—yes, it opens the Firefox extensions manager :-)

Another nice feature is the recent

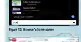

Figure 13: Browser's home screen

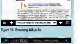

Figure 14: Browsing Wikipedia

Figure 15: Watching an YouTube video

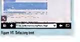

Figure 16: Selecting text

history view. Activate it by quickly swapping your finger in from outside the right-hand side of the screen, or tapping the icon that shows an arrow pointing toward left, next to the full screen icon on the right bottom portion of the screen. History shows you all the pages you've visited in the past before striking on the current page. You swap your finger to the left to go back and open a page from history.

All in all, using this browser is a pleasant experience—who knew browsing on a mobile device could be so interactive? And thanks to the physical QWERTY keyboard—which although needs a bit of getting used to—you actually feel that size doesn't really matter that much.

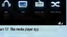
Figure 17: The media player app

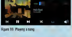
Figure 18: Playing a song

Figure 19: Exporting a photo to Facebook

Figure 20: Tagging an image

Things multimedia

Let's start with the default media player (Figure 17). Now I had a few songs on the microSD card I use in my LG phone. So, I hot-swapped it and launched the media player. To my surprise, it auto-detected all the songs. As expected, it's smart enough to understand ID3 tags and can categorise songs under albums, artists, or genres. Creating playlists is a no-brainer too. Besides, if the album art is inside the folder, the media player displays that too, in both the main media player (Figure 18), or the desktop widget. You can control the media playback using the desktop widget, too, and tapping the album-art area of the widget twice will launch the media player app for finer control.

The N900's stereo speakers have decent sound output, or you can use the headsets that come with the pack. The device is smart enough to maintain different volume levels between the headset and the device's stereo speakers. This is a nice feature considering, normally you'll have the volume at max when using the speakers, while if you plug in the headset you'll probably go deaf at a full volume output. However, when you plug in the headset, the volume switches to a much lower level, and you can increase (or decrease) the level that's most comforting to you.

Apart from the audio player, the media player is also home to Internet radio and videos. The radio has some preset stations—mostly European I guess. But adding new stations is as easy as accessing the add radio bookmark button, entering the station's name and the m3u stream URL. So, within seconds I was happily enjoying RadioParadise—my favourite station.

Watching videos is a pleasure, thanks to the high resolution and the wide-screen display. I copied a DivX file and the playback seemed crisp. This makes it a handy device when you're travelling and want to catch up on a movie to kill time. :) Also, while viewing a video, right from within the player, you have the option to share the video—over Bluetooth, as a mail attachment or via an online service. The online services that you can use are Flickr, Facebook, and OVI. Setting up these accounts is a one-time business, and next time onwards sharing is as simple as a few taps. To check the feature I shot a video using the phone's camera, and soon I had that posted on my Facebook page. A conversation followed on my Facebook page, as people couldn't make anything of that video—it was shot in a hurry :-)

Talking about the camera, it's a 5 megapixels backed by a double-LED flash. The picture quality is comparable to that of digital cameras—which makes the N900 one of the coolest GNU/Linux-powered one-size-fits all device. Click a picture, geo tag it, and share it instantly using the online services, if you have some sort of connectivity.

Back to the media player—it would have been really handy if it integrated the picture viewer application too. Nonetheless, it's available as a separate application, and has a decent set of features. Opening it gives you access to all the pictures/screenshots inside the device, where ever they are located—its guesswork is intelligent enough, just like that of the media players. Tapping an image opens it in full-screen mode as expected. To browse your collection you can swap your finger right/left to view the next/previous image—or better use the slideshow view. You can rotate the display to all the four sides of the screen by simply rotating the device. The accelerometer, however, didn't work while playing videos.

There are many other features and you should explore them yourself to see the magic of Maemo's multimedia capabilities. The best part is it even comes with an FM transmitter—so you can broadcast what you're listening to. I couldn't test how it works in the short span of time I had.

Conversations

Call, IM, SMS, e-mail contacts—that's conversation. Let's start with contacts. Open the keyboard and type a contact's name, and it automatically opens the list and starts filtering. Tap a contact, and if things are set up properly you'll have the option to call, SMS, Skype and IM (based on your IM accounts and settings). See Figure 21. Since my GTalk account is already set up, I have the options for GTalk call and chat, right there. Thanks to the VGA camera at the front, video conferencing over GTalk is a pleasure. Take a look at Figure 22.

Both SMSs and IMs figure in threaded lists for each contact I've had a chat with. For e-mails, you have to configure a separate mail client application—but it uses the same address book as far as integration is concerned. Talking about integration, you better have your contact list set up properly with the e-mail, IM, phone number etc. fields.

In my case, my phone numbers came from the SIM card while the IM contacts came from the GTalk account. So what I had were multiple entries for the same contact. A good feature that the contacts program provides is to manually merge a contact with another—and this is how I had managed to put things in place. An easy way would have been if only my Gmail contacts list had all the details filled in.

When you're busy with conversing with a contact, or just working on something else, a nice translucent yellow pop-up appears on the top of the screen to show a part of of the incoming SMS/IM/e-mail (Figure 23). It stays for a few seconds—you can immediately click to view it—and disappears in the background, leaving the dashboard icon glowing to notify you of a new message.

For normal calls, I found volume to be a bit low. Sometimes I really had to push the phone against my ear. A good thing is that the device comes with a proximity sensor. So the moment the phone is brought near the ear, the display shuts off. This is a very important feature; otherwise, imagine pressing the phone close to your ear and somehow hitting the call end button in the way.

Package manager

The App manager (Figures 25, 26 and 27) works just like Synaptic. After all, it's *apt-get* at the backend. By default, it comes with four repositories (package catalogues) enabled—Nokia Applications, Nokia System Software Updates, Ovi and maemo.org. The last one is where community supported apps are hosted, and this is the main resource for most of the apps you'll install to the device. The repositories sync with the online repos at regular intervals, and if there are updates to any application, you get a notification in your system tray just like you get in your desktop computer.

To begin with, I had some issues with the app manager. The official Nokia repos just won't sync, and because of this I was also unable to install a lot of apps—because either the app manager wasn't able to find the package, or some won't install as the packages they depended on were hosted in the Nokia repos. Perhaps, this is also the reason I wasn't notified of the Maemo update in the beginning that made me flash the device

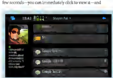

Figure 21 Addressbook of a contact

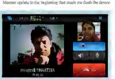

Figure 22 Video conferencing on GTalk

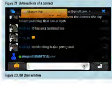

Figure 23. IM chat window

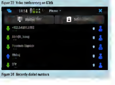

Figure 24 Recently dialled numbers

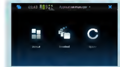

Figure 25: Applications manager

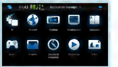

Figure 26: Apps are divided into categories

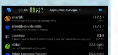

Figure 27: Peeking inside the 'Network' packages category

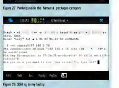

Figure 28: SSHing to my laptop

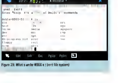

Figure 29: What's under N900's / (root file system)

manually. However even after that the problem remained. The solution was to disable the three Nokia repos and enable them again. This made them work. A weird sort of a bug I'd say.

The packages are listed under different categories (network, multimedia, graphics, etc) which makes it easier to find apps. You can also search apps by their names or even description. I installed AbiWord, SSH client and server, gPodder podcast manager, a few games like SuperTux, and a few more to check them out, and they all worked fine.

In fact, when I read that the KOffice viewer (aka FreOffice) alpha version was also released [*www.kdedevelopers.org/node/4143*], I couldn't wait but add the Maemo Extras-Devel repo and try it out. Although it's not recommended to install apps from this repo, don't we do it all the time on our desktops anyway? Besides, if something goes wrong, we can always flash the device because even that's so easy. By the way, this repo has lots of additional apps, so I highly recommend it.

Overall, you might not find as many apps for N900 as you would for an Android phone. But then, there are enough and more will be on their way as users join the community and help out in porting apps. Besides, now that the OVI Store for Maemo has also gone live, I'm sure commercial software writers will also take notice.

By the way you can also launch the X Term and install packages the *apt-get* way. You'll need to install the *rootsh* package. This helps you *su* to root (by running the *root* command), and you'll know what to do next. However according to reports, by default, you can't install Deb files using *dpkg*. For that you'd have to take the "Red Pill". Read more here: *http://nokiaaddict.com/2009/12/17/install-deb-files-on-the-n900-take-the-red-pill*

Isn't it amazing when a mobile devices gives you so much control? And this is all legal. No more do you need to google for "How to jailbreak N900" and stuff like that.

X-term and file system

Well, the longtime users swear by the shell, and the N900 comes with a busybox shell out of the box. You can SSH to other computer from the N900 or vice versa (Figure 28). Could things have been any easier? And switching to root is also simple as I explained earlier.

A thing to take note of is the fact that the keyboard doesnt have the tab and escape keys due to lack of space. However, these are vital when youre using a shell. I'm not surprised that the developers didnt fail to take notice of the fact and enabled these as part of X Terms UI

Using the root shell you can peep into the / directory to see whats in there (Figure 29). The N900 comes with a 32GB storage as we know Lets see how the device is partitioned after installing the *cfdisk* utility—take a look at Figure 30 The storage is divided into three partitions—approx 29GB vfat for your own storage: approx 2GB of ext3 root partition which is enough to install a lot of apps; and around 800MB of swap space (that adds up to the 256MB of RAM to give you a gigabyte of total memory) Of course, you can check the availability of space using a graphical tool from within the device, too (Figure 31).

Multitasking and performance

Well I shall conclude with this section Although I probably havent even covered 50 per cent of what the Maemo and the N900 combo offers, I ve run out of space and time :-(

Multitasking is a pleasure here. You can listen to your music, chat with multiple buddies over GTalk or SMS, browse the Web, work on a document using AbiWord, etc—all simultaneously. And you can switch between these apps seamlessly using the wonderful Dashboard application The 600Mhz OMAP processor is powerful enough to take care of this sort of a load However sometimes I did experience a lag in response time when too many apps were open—and I'm talking more than eight here. It makes me wonder what would happen if the next version of the device comes with 512 MB of RAM and a processor clock speed of 1GHz (to match the Nexus One)

On the down side, Id have preferred a battery with a higher capacity 1320mAh is just not enough for the kind of multitasking we can do on the N900, and I saw the battery ran out of power within 6-8 hours of continuous usage. Thankfully the USB cable doubles up as a charger when connected to a computer so you dont have to carry your charger separately everywhere you go.

Overall Id suggest the N900 to anyone who loves GNU/Linux. Android is not really for them—you cant really get into the system internals on that as easily as this one. The pricing and the launch dates in India are still a mystery But if we go by the European prices, an unlocked N900 should come for anything between Rs30,000-Rs35,000

Resources and useful links

- Nokia's Maemo website: http://maemo.nokia.com
- Maemo.org community: http://maemo.org
- Maemo.org Wiki: http://wiki.maemo.org/Main_Page
- Maemo.org forum: http://talk.maemo.org
- App reviews and how tos: www.maemo-guru.com

By Atanu Datta

He likes to head bang and play air guitar in his spare time. Oh, and he's also a part of the LFY Bureau

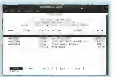

Figure 30: N900's partition table

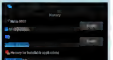

Figure 31 How much space is available

Figure 32 The Conky system monitor application

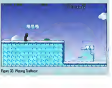

Figure 33 Playing TuxRacer

Starting Out with Maemo App Development

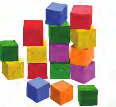

Maemo development can be divided into two broad categories — application development and platform development. In this article, we will not go into excessive details about developing the platform but focus more on getting started with application development on Maemo 5

Maemo is derived from the Debian GNU/Linux distribution and as a package it combines the Maemo OS and the SDK. Maemo consists mainly of the software stack from the Linux kernel to the Maemo APIs and the Hildon UI framework. Using Maemo, you can develop your own applications atop the default Hildon UI provided by the supplier

The UI architecture of Maemo is based on the GNOME desktop with many derivations from the GTK+ framework, GStreamer for multimedia, GConf for configuration and an XML library

Setting up the development environment

All the application development for Maemo is done with the help of the Maemo 5 SDK. The SDK creates a Maemo development environment on your system. It's a sandbox environment built on a tool called Scratchbox—it behaves as a Maemo based device on your system and has the tools to support the development. See the sidebox on Scratchbox for more information

It is similar to a hardware emulator designed as a software simulator. I will not go into the details about installing Maemo on your Linux system. If you use Debian derivatives then the chances are that a binary package is already available in your repositories, which you should search for and install. I still urge you to read through the official installation guide on the Maemo Community website at *http://wiki.maemo.org/Documentation/Maemo5_Final_Installation*. Once you have finished with the Installation of the SDK, you can begin tuning it for development.

For the emulator in the Scratchbox to work, you need the *binfmt_misc-feature* module loaded. Check whether or not it is loaded with *lsmod | grep binfmt*. If not, then as the root user execute *modprobe binfmt_misc*. If it does not work, check the documentation for your Linux system. It may be directly embedded in the kernel instead of being a module. In the worst-case scenario, you will have to compile your own kernel from source.

Once you develop the application, you will need an X server to test it. For this, you need to Install a pseudo X server. We will be focusing on Xephyr in this tutorial. Xephyr is an X server that can emulate 16 colour depth for its clients even while acting as a client to a 24-bit depth real X server. Xephyr also implements modern X protocol extensions. Install Xephyr from your distro's repositories.

As mentioned before, Maemo SDK uses Scratchbox as the cross compilation environment. Therefore, you must install Scratchbox first before attempting to install the SDK. The preferred way to install scratchbox and SDK is to use the automated installation scripts, though manual installation steps are also provided. Detailed instructions are found at *http://wiki.maemo.org/Documentation/Maemo_5_Final_SDK_Installation*. In short, the procedure for Debian-based systems is as follows.

- Install Maemo SDK build environment.
- Set up and create development targets.
- Start developing.

Another very good source of information on setting the development environment can be found at *http://maemo-sdk.garage.maemo.org/install_ubuntu.html*.

Getting started with the development

With the SDK installed and working in place, you can test the development environment. Create the usual sample command line-based Hello World application in C.

```
#include <stdio.h>
int main(int argc, char** argv)
{
   printf("Hello world\n");
   return 0;
}
```

Save the file as *sample.c* in the root of your development environment. Remember you should compile this program under the SDK and not your normal system. Unless you build the program under the SDK, you'll not know whether the SDK works or not.

Now test the application by issuing the following command.

```
[aditya@arctic ~] > gcc -Wall -g sample.c -o sample
[aditya@arctic ~] > ls sample*
sample* sample.c
```

GCC's -g flag is used to add debugging symbols to the output file. The *-Wall* flag will instruct the compiler to enable most of the syntax and other warnings generated on compilation. *-o sample* means the output binary filename is sample. Now run the program (*./sample*) and you shall see 'Hello World' printed on the terminal. Congratulations, it means your environment is up and running.

A step further: diving into the GUI

It's time now for some graphical sorcery. A simple GTK+ based window is demonstrated in the sample program below. *The program is taken from the official Maemo Documentation and is released by Nokia under the CC license.*

```
#include <stdlib.h> /* EXIT_* */
/* Introduce types and prototypes of GTK+ for the compiler */
#include <gtk/gtk.h>
int main(int argc, char** argv)
{
/* We'll have two references to two GTK+ widgets. */
GtkWindow* window;
GtkLabel* label;
/* Initialize the GTK+ library. */
gtk_init(&argc, &argv);

/* Create a window with window border width of 12 pixels and a
title text. */
window = g_object_new(GTK_TYPE_WINDOW,
"border-width", 12,
"title", "Hello GTK+",
NULL);

/* Create the label widget. */
label = g_object_new(GTK_TYPE_LABEL,
"label", "Hello World!",
NULL);
/* Pack the label into the window layout. */
gtk_container_add(GTK_CONTAINER(window), GTK_WIDGET(label));
```

Scratchbox

Maemo development is done in a Sandbox environment. Scratchbox delivers this sandbox environment on your system, providing development tools and at the same time, isolating the rest of the system from itself. Since there is a virtual environment available for development, cross-compiling is very easy.

The name 'Scratchbox' is a derivation from "Linux from Scratch" + "chroot jail" (sandbox). Scratchbox uses the Linux feature of changing the path of a file that a process can see. It changes the root (/) of the system to its own directory (same thing done in the case of a jail shell). This is the isolated area where all the work is done—a sandbox. Apart from directory switching, diversion of system calls, library references and other small changes makes development easier.

So, basically the Scratchbox is:

- a process of implementation of an isolated sandbox environment.
- a collection of tools to make cross compilation and development easier.
- a collaborative environment supporting multiple developers with individual profile supports.
- a mechanism called *sbrsh* that supports execution of the binaries directly on the target system.
- a cross-execution simulator generally implemented through QEMU.

But Scratchbox does not end here. You can develop your own applications under the Scratchbox environment. Debian users will get a favour here as well. It includes some integration of Debian package management so that once the source files are set up correctly, all you need is to write a couple of configuration files and you can generate binary packages for various architectures. There are many other features integrated in Scratchbox and you can only enjoy it fully when you know them all.

```
[illegible]

/* Start the main event loop. */
g_print("main: calling gtk_main\n");
gtk_main();

/* [illegible] play a message to the sta[illegible] */
g_print("main: returned from gtk[illegible]
[illegible] n EXIT_SUCCESS;
```

Save the file, say *sampleGUI.c*. Now all you have to do is compile the program. Remember that you will need GTK+ on your system to be able to compile the program successfully, otherwise it will not work. There are some specific flags that need to be passed while compiling. The flags required can be known using the *pkg-config* tool. *pkg-config --list-all* | sort will give you the flags. Issue the following command to compile the program for now:

```
[aditya@archie ~]> gcc -Wall -g sampleGUI.c \
`pkg-config --cflags gtk+-2.0` -o sampleGUI
```

The detailed information about *pkg-config* help is available in the documentation of Maemo and is not a part of the discussion here. The *pkg-config* automation is possible, but, at this juncture, it is better to use manual flags. If some problems still persist, you might be missing some libraries that need to be explicitly linked. This can be sorted out using the *pkg-config* libraries:

```
[sbox-DIABLO_X86: ~] > gcc -Wall -g sampleGUI.c \
`pkg-config --cflags gtk+-2.0` -o sampleGUI \
`pkg-config --libs gtk+-2.0`
```

If you still encounter any compilation errors, the best way is to find someone on IRC channels and work these out. If all goes well, you can run the application using the following command:

```
[aditya@archie ~]> ./sampleGUI
```

It will not work unless you have started the virtual X server. Now to run the application on the actual environment emulator, run xephyr as:

```
[aditya@archie ~]> Xephyr :2 -host-cursor -screen 800x480x16 -dpi 96 -ac -kb &
```

Figure 1 shows the emulator start up. It should be outside the scratchbox environment. Back inside the scratchbox, give it explicit instructions to use the local socket and run GUI applications under the virtual server:

```
[aditya@archie ~]> export DISPLAY=:2
[aditya@archie ~]> ./gtk_helloworld-1
    main: calling gtk_main
```

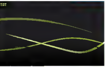

Figure 1: Emulator startup window

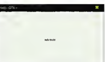

Figure 2: The 'Hello World!' dummy application

End the process with *Ctrl+c*. Now start the application framework from *af-sb-init.sh start*. It's a script to run the emulator environment. The screen similar to the one shown in Figure 2 will be presented to you.

What's good about Maemo

The best thing about Maemo is its open source software. More importantly, the entire documentation has been released under the Creative Commons license (you can find it on *www.maemo.org*).

In addition, the community has developed support for the Qt framework and it's also possible to develop applications in languages other than C, such as C++ or Python. With the added support to create a modular structure for the widgets, Maemo development of late has been wonderfully easy and very intuitive. And with more community involvement, it could perhaps evolve into the best mobile computing platform in the world.

By: Aditya Shevade

National Talent Scholar, Aditya Shevade, is a final year electronics engineering student, who takes keen interest in programming and electronic design. A Linux user for more than three years, he enjoys playing the keyboard and is a good photographer. To know more about him, visit www.adityashevade.com

CONTENTS

LINUX ForYou

FEBRUARY 2010 • Vol 07 No 12 • ISSN 0974 1054

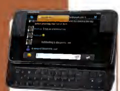

Introducing Maemo 5

Maemo is yet another open source mobile OS platform. We take a look at the Maemo powered N900 from its usability point of view to understand how it has matured over the years We also tell you how easy it is to start writing applications for Maemo.

38

FOR YOU & ME

ADMIN

GEEKS

Opening Up with Google

Although Google is the most popular Web-based search engine used by everybody not many know that it's one of the biggest supporters of open source. For more on Google's association with open source, here's the story

In January 1996, Larry Page and Sergey Brin, both students at Stanford University, began working on a search engine which came to be called Google and took the world by storm. Initially Page and Brin extensively used open source software such as Linux and GCC internally. One of the reasons for using open source was, perhaps, the cost savings. However, the power that open source offered them—the ability to study and customise the source code to their preference—played a bigger role.

The creators of Google used free software wherever possible, even for corporate processes like accounting, record-keeping, etc. The use of free software for Google's systems continues till date, with nearly all of its servers running on Linux. Page and Brin also used free software libraries such as OpenSSL, zlib, PCRE and MySQL. Since Google depended extensively on free software, its creators felt obliged to contribute to the various open source projects that had led to Google's success.

If the software improves, everyone benefits, including Google. So why keep the in-house fixes and improvements to the code a secret?

Helping open source projects

Google has always tried to maintain a good relationship with various open source projects. It hired Chris DiBona as the manager of its open source programs. DiBona oversees licence compliance and supports the open source developer community through programs like Google Summer of Code, and through the release of open source software projects and patches on Google Code.

Google has a separate blog on open source: the Google open source blog, which carries updates on all open source-related projects and programs that it is involved in.

Open source summits and conferences like the Ubuntu Developer Summit (November 2006 and December 2008) have even been held at the Google campus. Several of its employees work on open source projects. Google has also hired many known free software developers like Bram Moolenaar (Vim), Jeremy Allison (Samba) and Andrew Morton (Linux kernel). These developers work on their pet software while Google pays them.

Google also financially supports a number of open source products. It is one of the main sponsors of Firefox. It is because of this that Firefox has Google as the first entry in its search and 'awesome' bars.

Google employees have also made significant contributions to the Wine project. They played a major role in making Adobe PhotoShop CS2 work on Linux via Wine. This alone has resulted in many more people using Linux as their primary OS.

Google helped another open source project—Hunspell, which is used for spell checks in several applications including OpenOffice.org, Thunderbird, Firefox, and Mac OSX (10.6 onwards). The search giant subsequently decided to use Hunspell in its Chrome browser. When Google wanted to do something more, it scanned its huge database to find the 1,000 most commonly occurring words in different languages. It then compared this list with Hunspell's list of words and added those that were missing. Certainly an innovative way of helping Hunspell, something that Google is famous for.

Google's Linux software

Google releases Linux versions of its own software as much as possible and has released some very useful software for Linux—Picasa, Google Earth, etc. This allows its programs to be used by a wider base of more technically inclined users. It is unfortunate, though, that the latest version of Picasa has not been released for Linux. This, according to Picasa, is because the previous versions were not used by as many people as they had hoped.

Google has also open sourced its programs to gain users' trust. Google Update, a program for Windows that keeps all the installed Google programs updated, became troublesome since it kept running in the background. To tackle the problem, Google decided to open source the software so that users could work without having to worry about a 'suspicious' program running in the background. Now Google Update can be used by other projects as well to update software. It is because of such initiatives that the search giant continues to gain the trust of people all over the world.

In 2007, Google bought Jaiku, a social networking and microblogging site. After two years, it found that it could not put much of its resources into the project. Instead of letting the project rot, Google decided to open source most of its code. This way those interested could develop it and use it in some other project.

Open source developer tools

Google Code provides APIs and developer tools for projects that Google is involved in, like Wave, Android, OpenSocial and many more. Google Code also provides a code-hosting service with features such as revision control, issue tracker and a wiki. All these features are free for open source projects, enabling developers to host their projects for free. The project details are linked to the developers' Google accounts.

The Google Web Toolkit (GWT) is a set of open source tools that allow Web developers to create JavaScript front-end applications. A version of GWT 2.0, with significant improvements and enhanced features, was released recently.

In just the last few months, Google has made a number of announcements, directly or indirectly affecting open source development. It has developed a programming language called Go, aiming to make application development easier and faster. It's the TIOBE Programming Language of 2009 [*www.tiobe.com/index.php/content/paperinfo/tpci/index.html*][1].

The SPDY protocol transports Web content faster, hence reducing the page load time. Google also open sourced many of its own JavaScript application development tools recently.

Encouraging student developers

Google organises and manages two very useful open source programmes to encourage students to contribute to open source. These are the Google Summer of Code (GSoC) for university students and the Google Highly Open Participation Contest for high school students. The idea to launch these programmes came from the founders themselves. These programmes help and encourage students to work on an open source project for a few months. Apart from earning some pocket money from the summer job, the students gain from having worked with Google.

In 2009, India had the second highest number of GSoC participants—the proposals of 101 students were accepted.

Cloud computing benefit

Perhaps the best way that Google has supported open source is by helping people move to Web applications and cloud computing. With applications like Google Mail, Google Calendar, Google Docs around, it really doesn't matter what operating system you use.

You may tell me that none of these are really open source. So what? If Linux does ever achieve a majority market share, do you think it won't have popular proprietary applications? At least most of Google's apps can be used via open source applications like Evolution, Thunderbird, Pidgin, etc. As these applications are all cross platform, it makes switching from Windows to Linux a simple task. And since Google makes an effort to ensure all the data stored in its servers is easily accessible, there isn't any fear of data lock-in.

On May 27, 2009, when Google Wave—described by the company as a 'personal communication and collaboration tool'—was launched, everyone knew it was a potential game changer. The best thing about Google Wave is that Google plans to open source all the code. Hence, unlike Gmail, you can have your own Wave services on your own servers. So those who don't like using Gmail because of either privacy issues or its non-free nature, can use Wave without worries. Extensions can be easily developed to add new features to Wave. Also, since Wave is a tool for real-time communication and collaboration, it will help open source developers communicate much better.

And after the platform matures, developers can leave IRC channels and mailing lists, and switch to Wave. Some artists and documentation writers are already using Wave to collaborate better.

The Android revolution

Of all open source software developed by Google, the one that has interested companies and individuals alike is the Android operating system. When Apple launched iPhone, users began to understand that the magic ingredient of a successful phone is its software. The iPhone helped create a demand for smart phones, and this meant other cellphone manufacturers had to look for a really good operating system for their phones.

Among Linux-based smart phone operating systems, Android is clearly the most successful. Google involved hardware developers in the creation of the Android, which was a good strategy. The Android has showcased Linux's scalability and adaptability in smaller form-factors. Android has also pushed the use of open source software in cellphones and helped put Linux in the hands of thousands who earlier depended on proprietary software.

Chrome and the Chrome OS

The Chrome browser has managed to gain more than 4.5 per cent of the market to become the No. 3 browser in the world. But the real change that Chrome engineered was that it forced its competitors to produce better products. Chrome also brought with it an innovative and completely unique interface. Firefox 4 plans to include some of the ideas from Chrome's interface.

Google also developed a new algorithm called the Courgette, which helps provide faster and smaller updates for Chrome. Since this was also open sourced, many other projects considered using it for their update system. It is unfortunate, though, that Google was recently sued for patent infringement over Courgette, and its use is somewhat uncertain now. This brings us to the point of why Google is interested in creating its own open source operating system, namely, the Chrome OS.

Google's primary revenue channel is online advertising. It accounts for over 90 per cent of Google's income. For an organisation so heavily dependent on people clicking on Internet ads, it is natural for Google to try to get as many people online as possible. The two main factors that stand in the way are hardware and software. The hardware market is as competitive as it gets, with razor thin margins and products available to suit all kinds of budgets. But it is the software market which, because of the monopoly of a few players, becomes the bottleneck. One can easily see this reflected in the price of netbooks. If a netbook is available for Rs 12,000, then does it really make sense to pay a few more thousands for a copy of Windows?

This is precisely why Google decided to make its own OS and give it away for free. Many worry that this will only add to the problem of fragmentation in the open source ecosystem. But then, there are many reasons why they need not be afraid of the Chrome OS. The Chrome will be an OS on a secondary system. It will have only Web applications, mainly for netbooks. Another reason is that the features on the Chrome OS will be available through the Chrome browser. This means that whatever improvements Google makes in its Web applications or its OS, will be available to you. And even if a certain feature is not made available, you can easily port it, as it is all open source.

The Chrome OS brings with it some other advantages as well. OEMs like Dell and HP would feel more comfortable with an OS backed by Google, which can match Microsoft and Apple in popularity. Individuals, too, would rather try out an OS from Google than something not known to them. And in the end, since Chrome is all open source, and uses the Linux kernel, it will help improve other free operating systems as well.

By: Praveen Thirukonda

The author is a third year computer engineering student who loves playing around with FLOSS. He is a huge sci-fi movie buff and also enjoys playing badminton.

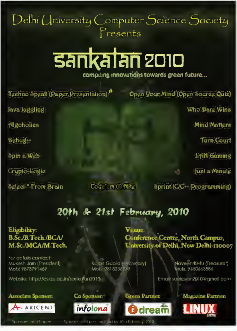

Delhi University Computer Science Society
Presents

sankalan 2010

computing innovations towards green future...

Techno Speak (Paper Presentation) #
Java Juggling
Algoholics
Debug++
Spin a Web
Crypto-Logic
Select * From Brain
Code'em @ Nite
Open Your Mind (Open Source Quiz)
Who Dare Wins
Mind Matters
Turn Court
LAN Gaming
Just a Minute
Sprint (C/C++ Programming)

20th & 21st February, 2010

Eligibility:
B.Sc./B.Tech./BCA/
M.Sc./MCA/M.Tech.

Venue:
Conference Centre, North Campus,
University of Delhi, New Delhi-110007

For details contact:

Mukesh Jain (President)
Mob: 9873791467

Rajan Gupta (Secretary)
Mob: 9818236739

Naveen Ketu (Treasurer)
Mob. 9650663586

Website: http://cs.du.ac.in/sankalan2010

Email: sankalan2010@gmail.com

Associate Sponsor: ARICENT

Co Sponsor: infolona

Green Partner: idream

Magazine Partner: LINUX

Sponsorships still open   # Synopsis must be submitted by 10th February, 2010

Blast
From the Past!

Automated snapshotting and back-ups for home users

How many times have you lost your data? Ever frantically searched for the video clip of that trip to the hills–only to realise it's lost, or rather destroyed by a rogue hard disk? Or realised, too late in the boardroom that the presentation you'd worked on last night to be made to your APAC head (that magically turned massive losses into sweet profits) just won't open? Instead, an Invalid Format, Corrupt File error shows up.

It happens to everybody. I used to be aloof to back-ups, until one day I lost my entire hard disk to a head crash. That was just 80 gigabytes of data, but now handling half-a-terabyte means that if this happens again, my life is more or less over. It will perhaps take me years to recreate that data. Besides, some of it is historic and is not found anywhere online. I'm speaking of archived articles, passwords and hexadecimal passkeys for websites, networks, digitised cassettes with the originals destroyed, digitised LP Vinyl Records with originals destroyed since I've no long-playing record player... you get it. Lots of data!

Some of us religiously keep back-ups of everything including the kitchen sink. But most of us are put off by the idea of backing up our data every once in a while. It's really boring, isn't it? And yes, automated back-up solutions *do* exist, but they are notoriously complex to set up. One does get scared after reading about some guy who spent days trying to set up his automated back-up software in Linux and who then went ahead and took a back-up. Yet when his disk crashed, he could not recover his back-up due to the complexity of the software!

Many companies have been trying to

address this issue. Perhaps most novel in its approach was Apple Inc when it created its Time Machine. This had no user-configurable options because there didn't need to be any. All Time Machine ever did was incrementally back up all of the hard disk to a Time Capsule (which was nothing but a NAS machine, or for practical purposes, a portable hard drive attached using the network instead of the USB). You could then click a button on the file manager and be shown a 3D screen where you would be able to recover the file from any previous date. Say a document was created on Day 1, 25 per cent of it was completed on Day 2, 50 per cent on Day 3, 15 per cent on Day 4 and it got corrupted on Day 5. Using Time Machine, you could then see the file as it existed on Days 1 through 4, and then recover any version of the file you desired, such as the 25 per cent from Day 2 or the 90 per cent from Day 4. It still exists, and is one of the many crowning features of the Mac OS.

However, it's the open world that offers the best solution available. OpenSolaris uses the ZFS file system and snapshots are built in as a feature. On OpenSolaris, it is just a matter of opening up Nautilus, right clicking the file, selecting the appropriate option and moving the Time Slider to the time you want. It doesn't get better than this.

Sadly, Linux doesn't get to use ZFS because of licensing restrictions, so we have to settle for Time Machine-like software. Presenting 'Back In Time', the king of automated snapshots and back ups, in Linux!

Time travel, anyone?

There have been many such software, most famously Time Vault and FlyBack, but none have ever got the equation right. Either the software is too complex to use, or it just doesn't work well. We've been neglecting the UNIX philosophy here, which is to do one task and do it well? But with *Back In Time*, the *attyachar* [harassment] is over.

Did I sound philosophical or boring? Well, back-ups *are* important and good back up software *are* hard to come by. So let's get on with the software.

Back In Time is basically a pretty GUI tucked on top of the following command-line utilities: *cp*, *rsync* and *ln*. Yes, it basically copies, links and syncs files—a very simple idea, but very effective.

For the first back up, or snapshot, all the specified data to back up is backed up into a unique time-named subdirectory of a master directory that you specify. For every subsequent back-up, *rsync* only copies the changed files to a new time-named directory, and *ln* creates links to the non-changed files from the previous snapshot directory. Pretty right?

It also means that you can browse your back-ups like a normal filesystem and don't need BIT (*Back In Time*) or third party software to restore a back up if your OS is trashed; you can manually do so by simply copying the files, if need be. It also means that because the snapshots are not compressed, the back-ups are larger than the size of your source. Sure, that's a negative point, but most of us keep our back ups on very large portable hard disks, anyway.

Figure 1: Back In Time user interface

BIT doesn't even have a daemon of its own. All the actual back ups are done by—wait for it—*cron*! Yes, that's it! Rather than reinvent the wheel, where you are scheduling an *rsync* job, all you do is add a line to the *crontab* and you are good to *go*!

Installing it

It's too simple for words if you're using Ubuntu, Fedora, Mandriva, or openSUSE. It's already in the Ubuntu, Fedora and Mandriva repos, so search for *backintime-common* and *backintime-gnome* for GTK and *backintime-kde* for KDE4.1. In openSUSE, you'll get it in the third-party packman repo, and in Ubuntu use the default package manager of your distro to install the package(s).

I've been using Slackware, and it's one of the few distros that need compilation from source code (I guess openSUSE has a build-service repo too). Added to that, I've been using Xfce4 rather than KDE4, and we have a unique installation scenario here.

The first run

You should only need to run this program once. Set it up, schedule the snapshotting interval, and then remove it. Yes, remove it. Because once the scheduling is complete, you don't need BIT any more, unless you want to change the snapshotting settings or need a pretty GUI to recover your files. Anyway, the first time you run it (be sure to run it as the root), the settings dialogue shows up (Figure 1).

The first tab is the *General* tab. It lets you select a place to keep your snapshots, and lets you specify a scheduling interval. I'd say the one-hour scheduling interval is just right: enough snapshots and not too many unnecessary back ups. Interestingly, there's a five-

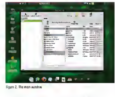

Figure 2. The main window

minute interval too for the paranoid ones I guess.

Next is the *Include* tab If you want a back-up of your entire system—excluding mounted drives, CDs and USB or network drives—youd add only one entry here- / (i.e. slash)

Move on to the next tab and configure the excludes here There are a few predefined patterns here to make sure temporary files and back ups arent backed up. Add a few folders to it. */media*, */mnt*, */proc* and */tmp*.

The next tab is *Auto remove* that lets you set which snapshots to remove automatically I think it s safe to leave all *options* enabled here. This optimises disk space usage

The Options tab has only one option to enable notifications and it should be enabled (duh!) In the next tab (*Expert Options*), it s a safe bet to keep the second option (i.e. run nice as a *cron* job) enabled

The main window

The main window is a rather clean design It helps you get your job done fast and do away with it (Figure 2)

On the top is a toolbar (no menu bar) with just the number of buttons required *Backup Now* (to take a manual snapshot) *Update Snapshots* (to rescan for snapshots) *Snapshot Name* (to specify a name for a certain snapshot which makes it really special). *Remove Snapshot*, *Settings*, *About*, *Help* and *Exit* Note that the UI is the same for both KDE4 and GTK+ they only look different.

Taking a back up is as simple as clicking *Backup Now* All it takes is a few minutes while *diff* is run to check for changes changed files are copied over and unchanged files are linked Restoring is even more of a no brainer—on the *Snapshots* pane click on a snapshot that you want to restore If you want to restore the whole back up click on *Restore*. If you want to extract just a few files, select your snapshot, then browse to the desired files on the rightmost pane and just drag them out to your desktop or some folder It's that simple!

Backing up to network servers

This really isn't rocket science Remember FreeNAS? You can just set up a NFS server like that (or set up an NFS on your back up server if you have one) and mount it on your PC to back up the data to the NFS share.

My thoughts

If I had a large external hard drive dedicated to taking back-ups—500GB drives dont cost much—I would have installed this on all of my distros. As it is I'm not averse to back ups I just wish I had disk space for them

Oh and forget the floppies tape drives and back ups on DVDs. No one takes them any more

By: Boudhayan Gupta

The author is a 15 year-old student studying in Class 9 He is a logician (as opposed to a magician) a great supporter of Free Software and loves hacking Linux. Other than that, he is an experienced programmer in BASIC and can also program in C++ Python and Assembly (NASM Syntax)

Score a Goal With Postgres!

Shrinking budgets and performance hungry business applications are making organisations sway the open source way when choosing database management systems EnterpriseDB the enterprise Postgres company is scoring consecutive goals on reliability performance scalability and cost comparisons against both its open source and proprietary opponents.

Most football freaks would agree on José Mourinho being the iconic coach of this decade, having consistently shown his ability at taking strong raw talent and transforming it into championship-calibre teams After ensuring FC Porto was decorated with medals of the Portuguese Liga, Cup of Portugal UEFA Cup and UEFA Champions League, he went on to guide Chelsea through two consecutive Premier League titles and helped inter win the Supercoppa Italiana and National Serie A titles In 2004 EnterpriseDB a solution provider of an open source relational database management system performed similar magic with PostgreSQL, a leading open source enterprise-class relational

> "While the total cost of ownership of a 16-CPU (four server, four CPUs/server) Oracle configuration is $1,261,600 over three years, the TCO of a comparably configured Postges Plus Advanced Server deployment is $215,760. This would mean a TCO savings of $1,045,840, or 83 per cent."
>
> **—Ashish Mehra, director India operations, EnterpriseDB Software India**

database management system (RDBMS). PostgreSQL was a promising star in the database market when it was originally built twenty years ago; EnterpriseDB helped transform the world class technology of PostgreSQL into an enterprise class product with the release of Postgres Plus Standard Server and Postgres Plus Advanced Server.

"Postgres Plus is an open source database typically deployed in high profile, mission critical applications. Our open source products provide capabilities that address a range of enterprise requirements (application development, enhanced performance, runtime management, usability and scalability), which enable organisations to build and deploy applications that solve critical business problems. We also provide 24x7 'follow the sun' technical support, training and professional services needed to support all phases of evaluation, development, deployment, and ongoing production of mission critical systems," says Ashish Mehra, director India operations, EnterpriseDB Software India.

Based on open source PostgreSQL, which is developed by the largest independent RDBMS community in the world, EnterpriseDB's Postgres Plus Standard Server and Postgres Plus Advanced Server are products suited for transaction-intensive and mixed load applications. Typically, Postgres Plus is deployed as an online transaction processing (OLTP) database in support of enterprise applications such as enterprise resource planning (ERM), customer relationship management (CRM), supply chain management (SCM) and other applications that require databases to support high concurrency, performance, scalability and security.

But with multiple players (both proprietary and open source) in the RDBMS market, what makes Postgres Plus products a cut above the rest?

Easy on the pocket

Being open source software, Postgres Plus is cost effective as compared to its strong proprietary counterparts like Oracle, Microsoft, and IBM. Postgres Plus Advanced Server includes a breakthrough suite of compatibility technologies that run Oracle applications with little or no changes at a much lower cost. "The most significant difference between the two databases is around cost. While the total cost of ownership of a 16-CPU (four server, four CPUs/server) Oracle configuration is $1,261,600 over three years, the TCO of a comparably configured Postges Plus Advanced Server deployment is $215,760. This would mean a TCO savings of $1,045,840, or 83 per cent," says Mehra.

NTT, a Japan-headquartered telecommunications company couldn't agree more. They switched from commercial RDBMS products for their telecom business operation and back office systems to PostgreSQL and have availed significant cost benefits. "We originally estimated to save $20-30 million for five years by deploying PostgreSQL to in-house systems. But after adding Postgres Plus Advanced Server to our options, we now estimate the savings to double," testifies Takeshi Tachi, senior manager, NTT Open Source Software Center, Nippon Telegraph and Telephone Corporation.

InMobi Global Media Adnetwork, formerly mKhoj Solutions, that opted for EnterpriseDB in 2007, saw similar savings. "We chose Postgres Plus to solve our pain areas, namely scalability and support issues. Postgres Plus was one tenth the cost of other RDBMS. Further, we have seen huge cost reductions in our IT costs thereafter," says Mohit Saxena, VP (technology), InMobi Global Media Adnetwork.

Apart from the most obvious way of reducing IT expenditures by saving on commercial licensing fees, EnterpriseDB offers numerous strategies for organisations to cut costs. New line of business (LOB) application development and deployment, replication for reporting and business intelligence, and migrating non-mission critical and mission critical applications away from high-priced commercial databases are some leading strategies. "Query reporting and BI activities degrade the performance of OLTP applications. By offloading those queries to lower-cost non-production

CONTENTS

DEVELOPERS

COLUMNS

REGULAR FEATURES

LFY DVD: Sidux 2009-04 12 & Free BSD 8

Sidux is a desktop-oriented live distro based on Debian Sid. The latest version ships with Linux 2.6.32 and KDE 4.3.4. Install it to unleash the power of bleeding-edge software atop a Debian base.

FreeBSD is the most popular BSD OS around. Some of the highlights of v8.0 are: Xen Dom U; VirtualBox guest and host; hierarchical jails; ZFS – no longer in experimental status; continued SMP scalability improvements; revised network link layer subsystem; and a lot more.

LFY CD: Postgres Plus 8.4 & LFY 2009 Archive

Postgres Plus Standard Server contains all the new features in the community version of PostgreSQL 8.4 plus it has the most important enterprise add-on modules (e.g., database connectors, PostGIS geospatial data support, replication, connection pooling, distributed cache, and job scheduler), all pre-packaged and installed in a one-click installer.

Additionally find the complete *LINUX For You* magazine archive for the year 2009 inside this CD...

Note: All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 Unported License a month after the date of publication. Refer to http://creativecommons.org/licenses/by-sa/3.0 for a copy of the licence.

Editor
RAHUL CHOPRA

Editorial, Subscriptions & Advertising
DELHI (HQ)
D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 26810602, 26810603
Fax: 26817563
E-mail: info@efyindia.com

BANGALORE
No. 9, 17th Main,
1st Cross, HAL II Stage,
Indiranagar, Bangalore 560008
Ph: (080) 25260023; Fax: 25260394
E-mail: efyblr@efyindia.com

Customer Care
E-MAIL: support@efyindia.com

Back Issues
Kits 'n' Spares
D-88/5, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 26371661-2
E-mail: info@efyindia.com
Website: www.kitsnspares.com

Advertising
CHENNAI
Venkat Raman K.
EMS House, 28-A, Cathedral Garden Road
Near Palmgrove Hotel, Chennai 600034
Ph: 044-28275191; Mobile: 09444630572
E-mail: efychn@efyindia.com

KOLKATA
D.C. Mehra
Ph: (033) 22294788
Telefax: 22650094
E-mail: efycal@efyindia.com
Mobile: 09432422932

MUMBAI
Flory D'Souza
Ph: (022) 24950047, 24928520; Fax: 24954278
E-mail: efymum@efyindia.com

PUNE
Zakir Shaikh; Mobile: 09372407753
E-mail: efypune@efyindia.com

HYDERABAD
Syed Liyaqath Hussain; Ph: 09848333966
Adesh Kotak; Ph: 09703640547
E-mail: efyhyd@efyindia.com

SINGAPORE
Ms Peggy Thay, Publicitas Singapore Pte Ltd
72, Bendemeer Road, #03-26
The Luzerne, Singapore 339941
Phone: +65-6836 2272; Fax: +65-6297 7302
E-mail: pthay@publicitas.com,
singapore@publicitas.com

UNITED STATES
Ms Veronique Lamarque, E & Tech Media
Phone: +1 860 536 6677
E-mail: veroniquelamarque@gmail.com

CHINA
Ms Terry Qin, Power Pioneer Group Inc.
Room 419, Zhongfan Building, No. 402
Dong Sangxie Industry Park, No. 52 Zhuahua Road
Futian District, Shenzhen China, 518031
Phone: (86 755) 83729797; Fax: (86 21) 6455 2379
Mobile: (86) 13923802595, 18603055818
E-mail: terryqin@powerpioneergroup.com,
ppgterry@gmail.com

TAIWAN
Mr Leon Chen, J.K. Media
4F-7, No. 346, Sec. 2, Road Rd., Songshan Dist.,
Taipei City 105, Taiwan
Phone: 886-2-87726780 ext.10; Fax: 886-2-87726787

Exclusive News-stand Distributor (India)
INDIA BOOK HOUSE PVT LTD
Arch No. 30, below Mahalaxmi Bridge, Mahalaxmi,
Mumbai - 400034 Tel: 022- 40497401, 40497402,
40497474, 40497479, Fax: 40497434
E-mail: info@ibhworld.com

Printed, published and owned by Ramesh Chopra. Printed at Ratna Offset, C-101, DDA Shed, Okhla Industrial Area, Phase I, New Delhi 110020, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020.

> "We chose Postgres Plus to solve our pain areas, namely, scalability and support issues. Postgres Plus was one-tenth the cost of other RDBMS. Further, we have seen huge cost reductions in our IT costs thereafter"
>
> **—Mohit Saxena, VP (technology), InMobi Global Media Adnetwork**

databases. Postgres Plus' replication capabilities can provide real-time data warehousing at a fraction of the cost of Oracle's replication solutions. This eliminates OLTP performance problems while ensuring timely delivery of critical information to business stakeholders," says Mehra. For example, FTD, the worldwide leader in floral-related products and services, implemented the Postgres Plus Advanced Server Replication Server module to improve performance of their production systems while still meeting the needs of their vendor network. "By moving their vendor-facing order tracking system away from their Oracle-based production environment and onto replicated systems running on Postgres Plus Advanced Server, the performance of their production systems increased more than 400 per cent during the next peak ordering season – all while reducing the cost of operating this system by more than 80 per cent," says Mehra.

Rich features

The fundamental features of Postgres Plus Standard Server and Advanced Server distinguish them from their primary open source competitors like MySQL and Ingres. "PostgreSQL contains a single unified storage engine capable of performing extremely fast for all load types: OLTP (with full ACID support), reporting and mixed usage. MySQL, on the other hand, has pluggable storage engines specialised for particular types of usage. This type of configuration can result in bottlenecks as applications grow and change. Subqueries, too, are poorly optimised in MySQL. Further, user defined data types in Postgres Plus give users more options for customised solutions than MySQL, also building room for database enforced encryption," says Mehra.

The Oracle compatibility features found in Postgres Plus Advanced Server distinguish it from its distant cousin Ingres, another strong open source database competitor. "EnterpriseDB has added powerful and convenient Oracle features to Postgres Plus Advanced Server that are standard fare in the minds of many DBAs and developers, such as function packages, dynamic runtime instrumentation, query optimisation hints, Oracle SQL extensions, explicit transaction controls, and data dictionary views," says Mehra.

Further, Postgres Plus Standard Server and Postgres Plus Advanced Server come with many productivity tools like Postgres Studio (graphical tool for database and cluster creation/maintenance, SQL environment, and more), EDB SQL (SQL command line environment), EDB Loader (bulk data loader with error handling), DBA Management Server (for DBAs to handle monitoring, job scheduling, SQL terminal, software update management), DBA Monitoring Console (for resource usage and management), GRidSQL Monitoring Console for distributed data, and the Oracle Replication console.

Performance enhancer

EnterpriseDB's Postgres Plus products have multiple features across many facets of the database to help improve performance. While GridSQL partitions data across multiple machines and transparently performs queries, Asynchronous Pre Fetch can optimise regular index scans and bitmap index scans by issuing concurrent I/O requests to RAID (Redundant Array of Inexpensive Disks) hardware on Linux systems. On the other hand, Multi-Version Concurrency Control (MVCC) allows for high performance in applications that are both read and write intensive.

The best example of performance enhancement using Postgres Plus is hi5, a leading social networking site that deployed OLTP PostgreSQL installations, running on hundreds of servers. "The system supports the data transactions of more than 56 million active users each month. In June 2008, the company delivered more than 18.5 billion page views that were supported by PostgreSQL, serving nearly 11 million visitors to the site every day," boasts Mehra of Postgres Plus.

NTT, too, saw better performance when they deployed Postgres Plus for transactional telecom business applications. "We deployed PostgreSQL for years, but earlier versions resulted in performance degradation in long-time operations requiring periodic date-vacuum operation. After PostgreSQL 8.3 released in February 2009, the performance degradation had been completely resolved, and EnterpriseDB has contributed significantly to this improvement," says Tachi. Adding on, Saxena says: "Postgres Plus is the heart of our growth. With every release we have been setting new performance benchmarks. Postgres has helped us achieve even [previously thought] unachievable levels. This needs to be seen in the light of the fact that we are yet to use the capacity of Postgres Plus to its fullest."

Scalability

EnterpriseDB solutions are highly scalable. Postgres Plus Advanced Server with the GridSQL configuration allows a single database to be partitioned across multiple commodity hardware machines for expandability. "InfiniteCache is a Postgres Plus Advanced Server feature providing an infinitely expandable cache across commodity hardware providing flexible growth management. The compression technology allows multi-gigabyte databases to reside in memory for lightning fast performance," says Mehra.

Vouching for it, Saxena says: "InMobi is in a business where our traffic and our volumes are bound to grow exponentially. This is why we chose Postgres Plus and have faced no scalability issues so far. Today, we have two data centres in India and US. We plan to expand to Japan and see no problem with regard to managing databases so far."

Strong technical base

Postgres comes with strong technical base and support. The Postgres community provides robust and well-tested fixes and enhancements across multiple platforms. Postgres user communities (like User Groups) provide a wealth of collaborative information to users, ensuring vendor independence and a strong ecosystem. This is where Postgres surpasses competitors like Ingres with no community outside of the Ingres Corporation.

Foolproof support

EnterpriseDB's foolproof support has resolved the migration woes traditionally faced by organisations adopting open source RDBMS. EnterpriseDB offers a wide range of support options for Postgres — from free assets such as tutorials, product documentation and technical whitepapers to for-fee training, services consulting, and subscription offerings. "Our training classes are developed and run by Postgres community's leading contributors, like Bruce Momjian. Our Jump Start and Packaged Services offerings, like RemoteDBA and Architectural Health Checks, enable organisations to quickly develop skills and insight for fast realisation of the benefits of Postgres Plus," say Mehra. "Migration involves two things – customisation and support. Migration is very easy with EnterpriseDB as they have plug-ins for all our existing databases. Further, their support is unmatchable. We use remote support from them," testifies Saxena.

NTT, too, would vouch for it. Though their initial migration from commercial DBMS to PostgreSQL was a struggle, migration to Postgres Plus Advanced Server thereafter was a cakewalk. "Our estimates show that the total effort of [our] system migration project using Postgres Plus could be halved compared to the migration using the plain PostgreSQL," says Tachi. The smooth experience has made NTT Japan look at migrating legacy accounting systems from commercial OS and DBMS to Linux and Postgres Plus Advanced Server. "We estimated this migration to be troublesome and not compensate the licensing cost saving before, but now we feel it is worth to try by deploying Advanced Server (Postgres Plus)," says Tachi.

Road ahead!

Today, EnterpriseDB employs over 100 associates worldwide, supporting more than 300 customers. Its business partners include Red Hat, Synnex, Complere, Tomax, Contegix, Thump, Elastra, immixGroup, Fujitsu, IBM, Continuent, and many other IT bigwigs. With EnterpriseDB winning consecutive trophies for Postgres, we might see them perform the same magic that Red Hat did for Linux. History sure does repeat itself!

By: Vanisha Joseph

The author loves to experiment and writing for LINUX For You is her latest experiment. So, beware! Just a minute, she also happens to be a journalist during the day.

Getting Started with OpenLDAP

Part 1

Setting Up A Directory Server

Let us begin this series on LDAP by setting up openLDAP on our machines and unravelling its mysteries as the series progresses.

OpenLDAP is the open source implementation of the LDAP protocol to access a directory. Simply put, a directory is a repository of data, much like a database (but with significant differences) that is used to store huge amounts of data. Directories are optimised for fast search and retrieval. Directories are also fairly static—in the sense that the data in a directory is not updated very frequently.

LDAP or Lightweight Directory Access Protocol is a protocol used to access directories. OpenLDAP is an open source directory software suite conforming to the LDAP protocol (version 3) and it supports all major platforms. The current version of OpenLDAP is 2.4. There are significant changes in this version compared to v 2.3. In this article though, we will cover features that apply to the previous versions as well.

The OpenLDAP software suite consists of a directory server (*openldap-server*) and a client (*openldap clients*) to access the directory. To be able to query other LDAP servers, only the *openldap client* is required. The *openldap-client* package provides popular client tools such as *ldapsearch*, *ldapadd*, etc. The *openldap-server* software is required if you want to build a directory of your own.

On a Red Hat Enterprise Linux/Fedora system, the OpenLDAP RPMS are available as *openldap-server* and *openldap clients*. A simple check will reveal the RPMS installed in your system. To query the RPM database for the installed OpenLDAP RPMS, run the following command.

```
[vbg@vbg work-scratch]$ rpm -qa | grep openldap
```

In my system, the output was as follows.

Installing the OpenLDAP client

To be able to query other LDAP servers, you will need to install the *openldap-clients* package:

```
[root@vbg work]# yum install openldap-clients
```

With the client software installed, you can access other LDAP servers. There are quite a few LDAP servers on the Internet that allow anonymous querying, which implies querying a server without providing a user name and password, much like in anonymous FTP transactions where you can download files without providing a username/password for authentication.

To query a directory server conforming to the LDAP protocol, use the *ldapsearch* command.

A list of public LDAP servers can be found at: *www.keutel.de/directory/public_ldap_servers.html*. To test whether your OpenLDAP client software is installed correctly and is able to query remote LDAP servers, you may query any one of the listed servers. As an example, I chose to query the server x500.bund.de:

```
[vbg@vbg work scratch]$ ldapsearch -x -h x500.bund.de \
 -b o=bund,c=de
```

The -*h* option in the *ldapsearch* command arguments specifies the hostname to be queried. The -*b* option specifies the 'base' from which the search should begin and -*x* specifies 'simple authentication'.

Running the command displays a lot of output on the screen. Redirecting this output to a file will help us in understanding LDAP.

```
[vbg@vbg work scratch]$ ldapsearch -x -h x500.bund.de \
 -b "o=bund,c=de" > /tmp/ldapTGK1.txt
ldap_bind: Success (0)
        additional info: Bind succeeded
[vbg@vbg work scratch]$
```

The stderror output (which was not redirected by the above command) shows that binding (or authenticating and connecting) to the LDAP server was successful. Let us now look at the contents of this file, which contain the output of our query:

```
[vbg@vbg work scratch]$ less /tmp/ldapTGK1.txt
# extended LDIF
#
# LDAPv3
# base <o=bund,c=de> with scope subtree
# filter: (objectclass=*)
# requesting: ALL
```

The opening line shows us that the output is in the LDIF format. LDIF stands for LDAP Data Interchange Format. LDIF is used to represent LDAP data in text form. An understanding of LDIF is critical to be able to perform successful add/modify/delete/search operations on an LDAP server. LDIF files are useful for:

1. adding data to the LDAP directory,
2. modifying directory data,
3. deleting directory data, and
4. pretty much anything that has to do with visualisation of directory data.

The other lines also give us some more meaningful information. Note that all these are preceded with a '#'. This means that these are comments, not the results of our query. They just provide additional information. The other important information is the version of the LDAP protocol used for querying (v3), the base of the LDAP tree on which the search operation was performed, and filters applied to refine search results.

A detailed explanation of this is beyond the scope of this article. Maybe in a future article, we can delve deep into the workings of LDAP and uncover its hidden treasures.

Most directories do not allow anonymous writing. To be able to add/modify/delete data, you need to bind to the directory server host with a valid user name and password. Most directories also do not allow anonymous read (or search). Depending on your environment and sensitivity, you may also wish to do away with anonymous read on your directory. This will depend on your *ldap-server* configuration.

Configuring the LDAP client

If you query a single directory server (within your organisation), typing the hostname and search base on the command line each time might seem like a mundane and irritating task. You can customise the */etc/openldap/ldap.conf* (on Fedora/RHEL) file as shown below to specify an LDAP server and base. The OpenLDAP client utilities read the */etc/openldap/ldap.conf* file.

> **Warning:** Do not modify the */etc/ldap.conf* file. It is for a different purpose and is accessed by a different set of programs.

To be able to simply query the above server (*x500.bund.de*), just specify the following two lines in */etc/openldap/ldap.conf*:

```
HOST x500.bund.de
BASE o=bund,c=de
```

Make sure you comment out all the other lines at this stage. Once you have done this, just execute the *ldapsearch* command without the host and the base options:

```
[vbg@vbg work scratch]$ ldapsearch -x
```

This will connect to the server specified in the configuration file and will search from the base mentioned there.

LDAP is a TCP/IP protocol and an LDAP server binds to TCP port 389 by default. When using TLS, the default secure LDAP port is 636. Instead of specifying a host, an LDAP URI can also be used. LDAP URIs are of the form *ldap://host.fqdn*

The above LDAP search using LDAP URI instead of the hostname can be executed as follows:

```
[vbg@vbg-work scratch]$ ldapsearch -x -H ldap://x500.lund.de \
  -b [illegible] -de
```

...and the corresponding entry in */etc/openldap/ldap.conf* will be

```
URI ldap://x500.lund.de
```

Installing the OpenLDAP server

To query other directory servers is one thing and to be able to create your own directory server to store and serve data is quite another. The first step towards creating your own directory server is installing the required software. In Fedora, a *yum install openldap-servers* will do the job.

For the purpose of this article, I will create the *openldap-server* on a virtual RHEL 5.4 server and the client will reside on my virtual Fedora 12 host. The details of the test set-up are:

1. LDAP server hostname: vbg.knafl.org (192.168.122.1) – RHEL 5.4 (64 bit)
2. LDAP client hostname: vbg-work.knafl.org (192.168.122.244) – Fedora 12 (32 bit)

The LDAP server software installs the stand-alone LDAP daemon (*slapd*). The version of the OpenLDAP server on RHEL 5.4 is 2.3 (*openldap-servers-2.3.43-3.el5*). The *slapd* daemon binds on Port 389 and serves client requests based on its directory data. In Fedora/RHEL systems, the important locations are:

- */usr/sbin/slapd* – the stand-alone LDAP daemon binary program
- */var/lib/ldap/* – the folder containing directory data files
- */etc/openldap/slapd.conf* – the LDAP server configuration file
- /etc/init.d/ldap – init script for the LDAP daemon

It also creates the user and group LDAP. For security purposes, the *slapd* daemon runs with the privileges of this user.

```
[root@vbg ~]# cat /etc/passwd | grep ldap
ldap:x:55:55:LDAP [illegible]:/var/lib/ldap:/bin/false
```

Configuring the LDAP server

The main configuration file is */etc/openldap/slapd.conf*. For the init script to be able to read this file, the permissions of this file should be set to 640 and the owner group should be set to ldap:

```
[root@vbg openldap]# ls -l /etc/openldap/slapd.conf
-rw-r-----  1 root ldap 5731 Nov  6  2008 /etc/openldap/slapd.conf
```

Let's start configuring our very own LDAP server. As always, remember to keep a back-up of the main config file—just in case something goes wrong.

```
[root@vbg openldap]# cp slapd.conf slapd.conf.orig
[root@vbg openldap]# vim /etc/openldap/slapd.conf
```

The 'include' section of this config includes certain files from the schema directory. These files describe the structure of our directory. We will examine these in more detail in another article. For now, to get going quickly, the other important directives in the configuration file are:

```
database    bdb # this defines the database backend used by openldap
                # in this case - Berkeley Data Base Backend
suffix      dc=my-domain,dc=com  # This is the root of the Directory Tree
rootdn      cn=Manager,dc=my-domain,dc=com
              # This is the Distinguished Name of the superuser of the directory
# rootpw    secret
```

For a directory server, the DN and the base suffix are important concepts.

The base suffix of a directory specifies the top of the directory tree (much like "/" signifies the root of the Linux filesystem). All directory data is stored under this base suffix, much like folders and files in the traditional Linux filesystem. The DN (Distinguished Name) is the unique name of a directory entry (much like file names and directories—which need to be unique within the same parent folder).

LDAP is generally associated with DNS as most organisations would want their domain names mapped to their directory data. Therefore, if my domain name is *knafl.org*, this domain context can be represented in the directory as follows.

dc=org (tld)
dc=knafl (my domain name)

Therefore, in the directory the domain *knafl.org* can be represented in a hierarchical fashion as: dc=knafl, dc=org

We need to specify this base in our directory configuration file along with the superuser name and password.

The superuser is referred to as *rootdn* (the DN of the 'root' user). This should also fall under the base of the directory tree and could be referred to by the Common Name (cn) of *Manager*. Finally, we should set the password of the root user.

We can also set the rootdn password in clear text. This is a security risk and should be avoided. The password should be stored in an encrypted form. I am only including the 'text password' for demo purposes.

```
######################################
#Sample config file slapd.conf
database        bdb
suffix          "dc=knafl,dc=org"
rootdn          "cn=Manager,dc=knafl,dc=org"
                # Cleartext passwords, especially for the rootdn,
                # should be avoided.
                # See slappasswd(8) and slapd.conf(5) for details.
                # Use of strong authentication encouraged.
rootpw          FOSTKnafLinux
######################################
```

Apart from the above, I have made no changes to the default configuration file. Once we have set up the server, we need to start the LDAP daemon. The LDAP service will read the configuration file and start the *slapd* daemon.

```
[root@vbg openldap]# service ldap start
Checking configuration files for slapd:  bdb_db_open: Warning - No
DB_CONFIG file found in directory /var/lib/ldap: (2)
Expect poor performance for suffix dc=knafl,dc=org.
config file testing succeeded
                                            [  OK  ]
Starting slapd:                             [  OK  ]
[root@vbg openldap]#
```

To check whether our LDAP is up, we will query our server as follows:

```
[root@vbg openldap]# ldapsearch -x -h localhost -b dc=knafl,dc=org
# extended LDIF
#
# LDAPv3
# base <dc=knafl,dc=org> with scope subtree
# filter: (objectclass=*)
# requesting: ALL
#

# search result
# search: 2
# result: 32 No such object
#
# numResponses: 1
# [root@vbg openldap]#
```

We will receive the 'No such object' message from the server. This means that our LDAP server is running, but it could not locate our directory tree or the data.

Now we need to create the base data for our LDAP server. To do that, we need to create an LDIF file and use the client command *ldapadd*.

First, let us create a simple LDIF file for our directory base and save it in the file */tmp/ldap-base*:

```
[root@vbg ~]# cat /tmp/ldap-base
dn: dc=knafl,dc=org
dc: knafl
objectClass: top
objectClass: domain
```

Adding the entries to the directory:

```
[root@vbg ~]# ldapadd -x -D "cn=Manager,dc=knafl,dc=org" \
-W -f /tmp/ldap-base
Enter LDAP Password:
adding new entry "dc=knafl,dc=org"

[root@vbg ~]#
```

By default, we require rootdn privileges to write data in the directory. The *-D* option of the *ldapadd* command specifies the username (DN in LDAP terminology), and *-W* specifies that the password will be entered on the command line prompt. Finally the *-f* option specifies the name of the LDIF file that contains entries to be added to the directory.

Now when we query the directory, we get an output as follows:

```
[root@vbg ~]# ldapsearch -x -h localhost -b dc=knafl,dc=org
# extended LDIF
#
# LDAPv3
# base <dc=knafl,dc=org> with scope subtree
# filter: (objectclass=*)
# requesting: ALL
#
# knafl.org
dn: dc=knafl,dc=org
dc: knafl
objectClass: top
objectClass: domain
#
# search result
search: 2
result: 0 Success
# numResponses: 2
# numEntries: 1
```

Thus we have created the base of our directory. Now we need to add more data to it and also beef up the security of the LDAP server.

Let us first add some more data to create an address book. Let us add a section called addressbook in our LDAP tree under the base suffix. This section can be denoted as an Organisational Unit (OU). Therefore, our LDIF file will be:

```
[root@vbg ~]# cat /tmp/addressbook.base.ldif
dn: ou=addressbook,dc=knafl,dc=org
ou: addressbook
objectClass: top
objectClass: organizationalUnit
```

Let us add this OU to our directory:

```
[root@vbg ~]# ldapadd -x -D cn=Manager,dc=knafl,dc=org
-W -f /tmp/addressbook.base.ldif
Enter LDAP Password:
adding new entry "ou=addressbook,dc=knafl,dc=org"

[root@vbg ~]#
```

Now let us add some address book entries under the OU "addressbook". An example of an LDIF file with a user entry might look like the file displayed below:

```
[root@vbg ~]# cat /tmp/addressbook.ldif
dn: cn=Varad Gupta,ou=addressbook,dc=knafl,dc=org
cn: Varad Gupta
sn: Gupta
l: Gurgaon
street: M 37 Old DLF Colony Sector 14
st: Haryana
postalCode: 122003
homePhone: 0124 22222222
mobile: 09999999999
mail: varad.gupta@fosteringlinux.com
objectClass: top
objectClass: inetOrgPerson
[root@vbg ~]#
```

Add this entry to our directory:

```
[root@vbg ~]# ldapadd -x -D cn=Manager,dc=knafl,dc=org
-w FOSTERingLinux -f /tmp/addressbook.ldif
adding new entry "cn=Varad Gupta,ou=addressbook,dc=knafl,dc=org"
[root@vbg ~]#
```

Note the -w followed by the password of the rootdn instead of the -W used earlier. With the -w option, you can specify the password on the command line instead of being prompted for it.

You can similarly add other entries. Let us now query our directory from the Fedora client host.

```
[root@vbg-work ~]# ldapsearch -x -h 192.168.122.1 -b dc=knafl,dc=org
# extended LDIF
#
# LDAPv3
# base <dc=knafl,dc=org> with scope subtree
# filter: (objectclass=*)
# requesting: ALL
#

# knafl.org
dn: dc=knafl,dc=org
dc: knafl
objectClass: top
objectClass: domain

# addressbook, knafl.org
dn: ou=addressbook,dc=knafl,dc=org
ou: addressbook
objectClass: top
objectClass: organizationalUnit

# Varad Gupta, addressbook, knafl.org
dn: cn=Varad Gupta,ou=addressbook,dc=knafl,dc=org
cn: Varad Gupta
sn: Gupta
l: Gurgaon
street: M 37 Old DLF Colony Sector 14
st: Haryana
postalCode: 122003
homePhone: 0124 22222222
mobile: 09999999999
mail: varad.gupta@fosteringlinux.com
objectClass: top
objectClass: inetOrgPerson
#
# search result
search: 2
result: 0 Success
#
numResponses: 5
numEntries: 4
[root@vbg ~]#
```

In the next article, we will look at how to secure our basic LDAP server—adding an encrypted rootdn password, implementing access control, disallowing anonymous search, modifying and deleting entries, etc. If you have any queries, please feel free to mail me at varad.gupta@fosteringlinux.com

By: Varad Gupta

Varad is an open source enthusiast who strongly believes in the open source collaborative model not only for technology but also for business. India's first RHCSS (Red Hat Certified Security Specialist), he has been involved in spreading open source through Keen & Able Computers Pvt Ltd, an open source systems integration company, and FOSTERing Linux, a FOSS training, education and research training centre. The author can be contacted at varad.gupta@fosteringlinux.com

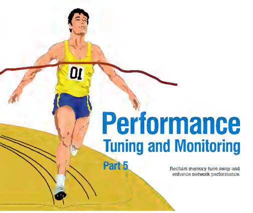

Performance
Tuning and Monitoring
Part 5

Reclaim memory, tune swap and enhance network performance.

Last month, we had worked on tuning our memory requirements—changing the kernel's behaviour for memory allocation, exploring the memory overcommit feature, and enhancing performance by tuning the ARP cache. To conclude this series on Performance Tuning and Monitoring, we will discuss memory reclamation, tuning swap and enhancing network performance.

States of virtual memory page

The operating system manages the virtual memory pages by reading a label known as a state. These labels (or states) help the virtual memory to take the required action on the page. So let's first understand the different states of a virtual memory page:

- *FREE* – the page is free or available for immediate allocation
- *INACTIVE CLEAN* – the page is not active and the contents are being written to the disk. There's nothing to be done on the page. In other words, the page is available for allocation
- *INACTIVE DIRTY* – dirty in this context means that the contents have been modified after reading from disk but not

written back on the disk, and the page is currently not in use
- *ACTIVE* – the page is currently in use by some process

Because memory is a limited resource, the kernel cannot store the dirty pages in RAM forever. As the kernel runs short of memory, the pages that have been dirtied, but are no longer used by a process, must be flushed back to the disk. This is handled by the *pdflush* kernel thread.

The default is a minimum of two threads. Use the following command to view the current status of *pdflush*:

```
[root@legacy ~]# sysctl -a | grep pdflush

vm.nr_pdflush_threads = 2

[root@legacy ~]
```

The idea behind *pdflush* is to avoid a sudden surprising load of heavy disk I/O. The goal is to write dirtied pages back to the disk at constant intervals. The default behaviour of *pdflush* is to write dirty pages to the disk as a kernel background task. The logic is simple: if the ratio of the number of pages that are dirty, to the total number of pages of the RAM available reaches a certain percentage, *pdflush* writes the dirty pages synchronously.

These are some of the other parameters that you can use to tune the behaviour of *pdflush*:

1. Tuning as per the length/size of memory
   - *vm.dirty_background_ratio* – sets the percentage (of total memory) of dirty pages at which *pdflush* starts writing to disk.
   - *vm.dirty_ratio* – sets the percentage (of total memory) of dirty pages at which the process itself starts writing dirty data to disk.
2. Tuning according to time
   - *vm.dirty_writeback_centisecs* – sets the interval between *pdflush* wake ups (in 1/100 of a second). Set zero to disable *pdflush*
   - *vm.dirty_expire_centisecs* – sets the time when the data is old enough (in 1/100th of a second) to be written to disk by *pdflush*

Here is a sample default configuration for *pdflush* values on my system:

```
[root@legacy ~]# sysctl -a | grep vm.dirty

vm.dirty_background_ratio = 10

vm.dirty_background_bytes = 0

vm.dirty_ratio = 20

vm.dirty_bytes = 0
vm.dirty_writeback_centisecs = 500
vm.dirty_expire_centisecs = 3000
[root@legacy ~]#
```

As you can see, *pdflush* will start writing the dirty pages to the disk as soon as 10 per cent of the total memory gets dirty. On the other hand, the process itself will start writing the data to disk as soon as 20 per cent of the total memory is dirty.

You can always reset these values by defining the new values in */etc/sysctl.conf* and then running sysctl -p.

Tuning swap

We all know what swap is and how it is important to us. We also know that swap is very slow compared to the physical memory (RAM). Naturally, when a system heavily relies on swap, it basically means we need more RAM.

A common controversial question is how much swap space to dedicate to a system. The answer depends on several factors. However, generally, you can use the following guidelines.

- If the amount of RAM is less than 1 GB, the ideal swap space should be twice the size of RAM
- If the amount of RAM is between 1-2 GB, the ideal swap space should be 1.5 times the size of

Editorial

Dear Readers,

"We seek your support."

That's the title of an advertisement on Page 36. This ad is for an event that's very dear to us—Open Source India, which some of you may remember as LinuxAsia. Allow me to share some of our thoughts that went into creating this ad.

There is no doubt that for any open source venture or event to be successful, it needs the involvement of various stakeholders. That's why we seek your support. We realise that in the open world, there are many different types of stakeholders involved, and when you combine the diverse forms of 'support' that each one can offer, we have no doubt in our minds that Open Source India 2010 can scale to being an event of global significance.

The shift in dates for the 2010 edition is an example of how your active participation makes things better. Thanks to a poll conducted amongst select readers and delegates of past events, we realised that September-October were preferred to our earlier practice of organising the event in February. Hence, we have shifted the dates to September 19-21, 2010.

Our shift to Chennai from Delhi was also a result of a similar feedback exercise. It was a move that catapulted Open Source India right to the top, with over 2,500 techies attending last year's show. The statement "The LINUX For You team adores Chennai" barely reflects the very positive sentiments we have for Chennai and the overwhelming support we received last year.

This year, we hope to do even better. Hence, we have started our preparations early. To get updates on Open Source India 2010, you can connect with the event through its website www.osidays.com, follow @osidays on Twitter or become a fan of its Facebook page.

> *This year, we hope to do even better. Hence, we have started our preparations early. To get updates on Open Source India 2010, you can connect with the event through its website www.osidays.com...*

The immediate support we seek: (a) open source experts who can share their knowledge as speakers, (b) individual enthusiasts who can join us as volunteers and help promote the event, (c) representatives of LUGs, colleges and IT organisations, who can partner with us and promote the event within their organisations, and everyone else who can be there, along with their friends and colleagues.

Chennai—we are getting ready for you. Hope you are ready for us!

Best Wishes!

Rahul Chopra

Rahul Chopra
Editor, LFY
rahul@efyindia.com

PS: Don't miss the article on Nokia's N900—a mobile phone that's powered by the Debian derivative OS Maemo. LFY team members swear by the coolness of this device, and some believe that the Linux community will prefer Maemo over Android!

RAM

- If the amount of RAM is between 2-8 GB, the ideal swap space should equal to the size of RAM
- If the amount of RAM is more than 8 GB, the ideal swap space should be 0.75 times the size of RAM

Let's now see how to improve the swap performance.

Ideally, dedicate a set of high performance disks spread across two or more controllers. If the swap resides on a busy disk, then, to reduce swap latency, it should be located as close to a busy partition (like */var*) as possible to reduce seek time for the drive head.

While using different partitions or hard disks of different speeds for swap, you can assign priorities to each partition so that the kernel will use the higher priority (actually high performance) disk first. You can set the priority using the *pri* option in your */etc/fstab* file. Here is an example of the */etc/fstab* file under such circumstances:

```
/dev/sda6        swap        swap        pri=4        0    0
/dev/sda8        swap        swap        pri=4        0    0
/dev/sdb3        swap        swap        pri=1        0    0
```

The kernel will first use the swap with a higher pri value—in our example, */dev/sda6* and */dev/sda8*. It will not use */dev/sdb3* until all of */dev/sda6* and */dev/sda8* are fully used. In addition, the kernel will distribute visit counts in a round robin style across all the devices with equal priorities.

The basics of network tuning

First, let's look at what happens when a packet is transmitted out of the NIC (network interface card).

1. Writes data to socket file (goes in transmit buffer).
2. Kernel encapsulates data into protocol data units.
3. PDUs are moved to the per-device transmit queue.
4. Driver copies PDU from head of queue to NIC.
5. NIC raises the interrupt when PDU is transmitted.

Now let's see what actually happens when a packet is received on NIC:

1. NIC receives frame and uses DMA to copy into receive buffer.
2. NIC raises CPU interrupt.
3. Kernel handles interrupt and schedules a *softirq*.
4. The *softirq* is handled to move packets up to the IP layer for the routing decision.
5. If it's a local delivery, then the packet is decapsulated and placed into the socket's receive buffer.

The 'txqueuelen' value, as reported by the *ifconfig* command, is the length of the transmit queue on the NIC. Packets are scheduled by a queuing discipline. I ran the ifconfig command to find out its value in my system:

```
[root@legacy ~]# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:15:C5:C0:38:76
[illegible]
[illegible]
[illegible]
          RX packets:27348 errors:0 dropped:[illegible] 0 frame:0
          TX packets:20180 errors:0 dropped:[illegible] 0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:22881585 (21.8 MiB)  TX bytes:3683788 (3.5 MiB)
          Interrupt:17
```

As you can see, it's set to 1000 by default. Let's now issue another interesting command related to the previous 'txqueuelen' value:

```
[root@legacy ~]# tc -s qdisc show dev eth0
qdisc pfifo_fast 0: root bands 3 priomap  1 2 2 2 1 2 0 0 1 1 1 1 1 1 1 1
 Sent 3601538 bytes 20149 pkt (dropped 0, overlimits 0 requeues 0)
 rate 0bit 0pps backlog 0b 0p requeues 0
```

The above command basically outputs the 'qdisc' for eth0, and also notifies us whether the connection is dropping packets. If it is, then we need to increase the length of the queue. Here's how we can increase the value to 10000:

```
[root@legacy ~]# ip link set eth0 txqueuelen 10000
```

We verify whether the new value has taken effect by running the *ifconfig* command again:

```
[root@legacy ~]# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:15:C5:C0:38:76
          inet addr:172.24.0.252  Bcast:172.24.255.255  Mask:255.255.0.0
          inet6 addr: fe80::215:c5ff:fec0:3876/64 Scope:Link
          UP BROADCAST MULTICAST  MTU:1500  Metric:1
          RX packets:27501 errors:0 dropped:0 overruns:0 frame:0
          TX packets:20208 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:10000
          RX bytes:22895522 (21.8 MiB)  TX bytes:3685984 (3.5 MiB)
          Interrupt:17
```

We can also control the size of the maximum send and receive transport socket buffers. We can use the following values in our */etc/sysctl.conf* file:

- *net.core.rmem_default* – the default buffer size used for the receive buffer

- *net.core.wmem_default* – the default buffer size used for the transmit buffer
- *net.core.rmem_max* – the maximum receive buffer used by an application
- *net.core.wmem_max* – the maximum send buffer used by an application

We can always find out the current values of these parameters using the *sysctl -a* command as follows:

```
[root@legacy ~]# sysctl -a | grep net.core.r

net.core.rmem_max = 131071

net.core.rmem_default = 114688

[root@legacy ~]#
[root@legacy ~]# sysctl -a | grep net.core.w

net.core.wmem_max = 131071

net.core.wmem_default = 114688

net.core.warnings = 1
[root@legacy ~]#
```

We'll now try to understand how to tune the size of the TCP buffer.

A high bandwidth connection is also known as a fat pipe, and a high latency connection is also referred to as a long pipe. For maximum throughput, we would like to transmit continuously. TCP buffers can be tuned by using the *net.ipv4.tcp_rmem* parameter for reads and *net.ipv4.tcp_wmem* parameter for writes. For each, *sysctl* holds three values: the minimum, the default and the maximum size.

For tuning the TCP buffer, we need to calculate the BDP (Bandwidth Delay Product), i.e., the length of pipe. The formula for calculating the BDP is: *Length of Pipe = Bandwidth x Delay (RTT)*, i.e., *L = B x D*. The value of BDP is in bytes.

Using *ping*, we can determine the RTT value. Let's assume that we are using a 100 MBps NIC, and the value of time that *ping* gives us is roughly 25.999 ms.

```
[illegible] / 8
327680 bytes
```

Let's look at the default values for the TCP read buffer in our system:

```
[root@legacy ~]# cat /proc/sys/net/ipv4/tcp_rmem

4096    87380       3366912
[root@legacy ~]#

[root@legacy ~]# sysctl -a | grep tcp_rmem
net.ipv4.tcp_rmem = 4096   87380   3366912

[root@legacy ~]#
```

Now I can change the values that have been calculated by BDP by using:

```
echo '327680 327680 327680' > \
/proc/sys/net/ipv4/tcp_rmem
```

...or I can edit the /etc/sysctl.conf file, as shown below:

```
net.ipv4.tcp_rmem = "327680 327680 327680"
```

Remember to run *sysctl -p* after making any changes to */etc/sysctl.conf*. Also remember to take a look at the man page for the *tc* command.

With this, I conclude this series of articles on 'Performance Tuning and Monitoring'. I do hope I have been able to share some good tips with you, and look forward to sharing with you something more exciting, soon. END

By: Alok Suvastava

The author is the founder of Network NUTS and holds MCP, MCSE, MCDBA, MCT, CCNA, CCNP, RHCE and RHCSS certifications. Under his leadership, Network NUTS has been a winner of the "Best Red Hat Training Partner in North India" for the last three years in a row. He has also been a trainer for the Indian Air Force, NIC, LIC, IFFCO, Wartsila India Ltd, the Government of Rajasthan, Bajaj Allianz, etc. You can reach him at alok at networknuts dot net.

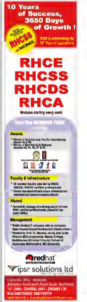

Part 8

The Art of Guard

Setting Security Contexts Using Policy Modules

Let's look at how to set default security contexts for files using policy modules. We will also use policy modules to set rules that will automatically transition the default type assigned to files and folders at the time of creation.

In the previous article, we had developed an SELinux policy module called test (version 1.1.1) for type enforcement. To do that, we had created a file called *test.te* which contained our allow rules and the new type(*lfy_t*) that we created in the system. We had written rules that allowed us to change the type of the */var/www/html/index.html* file from the default *http_sys_content_t* to *lfy_t*.

```
[root@vbg work]# ls -Z /var/www/html/index.html
-rw-r--r--  root root unconfined_u:object_r:lfy_t:s0  /var/www/
html/index.html
```

The only problem with this approach is that a restoration of file contexts would cause this file's security context to be relabelled back to *httpd_sys_content_t*. To check this, we can run the *restorecon* command.

```
[root@vbg work]# restorecon -v /var/www/html/index.html
restorecon reset /var/www/html/index.html context
 unconfined_u:object_r:lfy_t:s0->system_u:object_r:httpd_sys_
content_t:s0

[root@vbg work]# ls -Z /var/www/html/index.html
-rw-r--r--  root root system_u:object_r:httpd_sys_content_t:s0
/var/www/html/index.html
[root@vbg work]#
```

To avoid this issue and to set the default type of this file to *lfy_t*, we will need to add information to our policy module test. Entries that map files to desired security contexts are not specified in the *te* (type enforcement) files.

A closer look at our development folder reveals that the make command (used to compile our source to a binary policy module) has already created two new files: *test.fc* and *test.if*. We will understand the purpose and the utility of the *fc* (file contexts) file in this article. A sample *fc* file, *example.fc*, is provided by the SELinux development RPM. The default location of this file is */usr/share/selinux/devel/example.fc*. Let us have a look at the syntax.

```
[vbg@vbg work test selinux]$ cat example.fc
# myapp executable will have:
# # label: system_u:object_r:myapp_exec_t
# # MLS sensitivity: s0
# # MCS categories: <none>
#
/usr/sbin/myapp            --      gen_context(system_u:object_r:
myapp_exec_t,s0)
[root@vbg work]#
```

Simply put, this file has three columns:

1. The first column contains the name of the file (or a regular expression pattern) to be labelled (since everything is a file in Linux).
2. The second column indicates the type of the file (-- means any type of file).
3. The third column contains the default security context to be allotted to this file. It is very similar to file */etc/selinux/targeted/contexts/files/file_contexts* that we covered in the second article of this series. You may wish to revisit that article for the detailed explanation

To begin with, you can omit the use of macro *gen_context()*. We will cover macros in a later part of this series. After that, we will revisit the *gen_context()* macro.

For our current purpose, let's create the file, *test.fc*, containing the following line:

```
/var/www/html(/.*)?      system_u:object_r:lfy_t:s0
```

...and recompile our policy module. Do remember to increase the version number in the type enforcement file, *test.te*.

Let us update our test module:

```
[root@vbg work]# semodule -u /home/vbg/test/test.pp
```

To see the effect of our new *test.fc* file, let's check the current context of the file */var/www/html/index.html* and then restore it to its default context.

```
[root@vbg work]# ls -Z /var/www/html/index.html
-rw-r--r-- root root system_u:object_r:httpd_sys_content_t:s0 /var/www/html/index.html

[root@vbg work]# restorecon -v /var/www/html/index.html
restorecon reset /var/www/html/index.html context system_u:object_r:httpd_sys_content_t:s0->system_u:object_r:lfy_t:s0

[root@vbg work]# ls -Z /var/www/html/index.html
-rw-r--r-- root root system_u:object_r:lfy_t:s0      /var/www/html/index.html

[root@vbg work]#
```

Now we can see that the default context of the file has been set to *lfy_t*. What context does a newly created file get? Let us test this by creating a new file */var/www/html/test.html* by using the *touch* command.

```
[root@vbg work html]# touch /var/www/html/test.html
[root@vbg work html]# ls -Z /var/www/html/test.html
-rw-r--r-- root root unconfined_u:object_r:httpd_sys_content_t:s0 /var/www/html/test.html
[root@vbg work html]#
```

This shows that the default security context of the created file takes the type *httpd_sys_content_t*. Why is it so?

Let us create a new folder under */var/www/html/*:

```
[root@vbg work html]# mkdir /var/www/html/test
[root@vbg work html]# ls -Z /var/www/html
-rw-r--r-- root root system_u:object_r:lfy_t:s0      index.html
drwxr-xr-x root root unconfined_u:object_r:httpd_sys_content_t:s0 test
-rw-r--r-- root root unconfined_u:object_r:httpd_sys_content_t:s0 test.html
[root@vbg work html]#
```

The newly created folder also gets a security context with the type *httpd_sys_content_t*.

Now let us set the type of this new folder to *lfy_t* and then create new files under it.

```
[root@vbg work html]# chcon -t lfy_t /var/www/html/test
[root@vbg work html]# touch /var/www/html/test/test1.html
[root@vbg work html]# ls -Z /var/www/html/test/
-rw-r--r-- root root unconfined_u:object_r:lfy_t:s0  test1.html
[root@vbg work html]#
```

Surprise! The new files are of type *lfy_t*. This tells us about the default behaviour when applying security contexts to newly created files. As observed above, newly created files take the security context of the parent folder. Is it really so? Let us test this with a few more examples.

Security context of the root folder is *root_t*:

```
[root@vbg work html]# ls -Zd /
dr-xr-xr-x root root system_u:object_r:root_t:s0     /
```

Create a new file under the root folder and check its security context. Going by the above logic, the type of a newly created file (for example, *test.txt*) should be *root_t*.

```
[root@vbg work html]# touch /test.txt
[root@vbg work html]# ls -Z /test.txt
-rw-r--r-- root root unconfined_u:object_r:etc_runtime_t:s0 /test.txt
[root@vbg work html]#
```

What happened? The file *test.txt* under folder / (type – *root_t*) was supposed to bear a security context with type *root_t*. Instead, its type is *etc_runtime_t*.

Why is the type of the file not *root_t*. Why and how did the type of this file change? This is achieved through Type Transition. When a new object is created, we can create a rule in our policy module that can change the type of the newly created object. If no rule is present, the newly-created object inherits the type of its parents.

A new object is created by a process (a *subject*, in SELinux terminology). In the above example, when we created a new file */test.txt*, we did so by using the *touch* command. This was executed on the terminal windows (bash process). Let us check the type of the bash process.

```
[root@vbg work html]# ps auxZ | grep bash
unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023 1830 pts/0 Ss+    0:00
bash
unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023 2156 pts/1 Ss   0:00 bash
unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023 2182 pts/2 Ss   0:00 bash
unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023 2218 pts/3 Ss   0:00 bash
unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023 2244 pts/3 S    0:00  bash
unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023 2372 pts/4 Ss   0:00 bash
```

It does not matter which one of the above processes we used while creating the file, because all the processes have the type *unconfined_t*. Our policy could have had a rule which said:

"If a subject with type unconfined_t *creates a new object/subject whose the parent object/subject has type of* root_t, *set the new objects/subjects type to* etc_runtime_t."

In the absence of such a rule, the newly created object would simply have the type of its parent object. These are called type transition rules.

To check the number of type transition rules, in our policy use the *seinfo* tool:

```
[root@vbg work html]# seinfo
Statistics for policy file: /etc/selinux/targeted/policy/policy.24
Policy Version & Type: v.24 (binary, mls)

   Classes:           77    Permissions:       229
   Sensitivities:      1    Categories:       1024
   Types:           2804    Attributes:        226
   Users:              9    Roles:              12
   Booleans:         149    Cond. Expr.:       177
   Allow:         223880    Neverallow:          0
   Auditallow:        38    Dontaudit:      107109
   Type_trans:      9675    Type_change:        38
   Type_member:       97    Role allow:         18
   Role_trans:       184    Range_trans:      1602
   Constraints:       62    Validatetrans:       0
   Initial SIDs:      27    Fs_use:             22
   Genfscon:          82    Portcon:           403
   Netifcon:           0    Nodecon:             0
   Permissives:       36    Polcap:              2
```

Going by the above output, we see that there are 9,675 type transition rules in our default policy. To see all of them, simply run the following command:

```
[root@vbg work html]# sesearch -T
```

The output screen will display all the type transition rules. To see the top few type transition rules, we can use the command given below:

```
[root@vbg work html]# sesearch -T | head
Found 9720 semantic te rules:
   type_transition nx_server_t devpts_t : chr_file nx_server_devpts_t;
   type_transition user_wine_t fusermount_exec_t : process mount_t;
   type_transition staff_mozilla_t nsplugin_config_exec_t : process
nsplugin_config_t;
   type_transition user_mono_t fusermount_exec_t : process mount_t;
   type_transition xguest_t fusermount_exec_t : process mount_t;
   type_transition xguest_java_t fusermount_exec_t : process mount_t;
   type_transition ricci_modservice_t mdadm_exec_t : process mdadm_t;
   type_transition initrc_t ais_execserver_exec_t : process ais_execserver_t;
   type_transition hotplug_t dhcpc_exec_t : process dhcpc_t;
```

This gives us an opportunity to understand the syntax of a type transition rule, which can be specified as:

```
type_transition [Creating Subject type] [Parent Object Type] : [Object Class] [Type]
```

This helps us to understand the topmost rule as: If a new object of class *chr_file* (character file) is being created by a process with type *nx_server_t* under a parent object with type *devpts_t*, set the type of the new character file as *nx_server_devpts_t*.

Let us check our policy to see whether a type transition rule exists to change the type of a newly created file under / by the bash process.

As per the above syntax, it should be as follows:

```
type_transition        unconfined_t root_t        file              etc_runtime_t
(Creating subject: bash)  (Parent Dir: /)  (Class of the created object)  (New object type)
```

Let us search the policy for this rule. A rudimentary command for the same would be:

```
[root@vbg work html]# sesearch -T | grep unconfined_t | grep root_t | grep
etc_runtime_t
   type_transition unconfined_t root_t : file etc_runtime_t;
```

Therefore, we see that it is this rule that is changing the type of any file being created in the / directory. The above command also shows that this rule will not apply to any directories created under / since it is only applicable on the class "file". Let us check this out.

```
[root@vbg work html]# mkdir /test
[root@vbg work html]# ls -Zd /test
drwxr-xr-x  root root unconfined_u:object_r:default_t:s0 /test
```

Should the type of the folder */test* not be *root_t*? Then why is the type *default_t*? **Hint.** *Search for a type transition rule in the SELinux policies to apply to directories under the / directory.*

Let us now create a type transition rule in our policy module test so that any files created under the folder */var/www/html/* using the terminal (bash process) will be of the type *lfy_t*. The type transition rule will be:

```
type_transition unconfined_t httpd_sys_content_t : file lfy_t;
```

Let us add this line to our *test.te* file so that it now looks like (note the change in the version number) what's given below:

```
##########################
policy_module(test, 1.1.3)

require {
        type unconfined_t;
        type lfy_t;
        class file getattr;
        type httpd_sys_content_t; #you would need to add the required type as well


type lfy_t;
```

```
allow unconfined_t lfy_t:file { relabelto };
allow lfy_t fs_t:filesystem { associate };
allow unconfined_t lfy_t:file getattr;

###########Adding type transition rule ################3
type_transition unconfined_t httpd_sys_content_t : file lfy_t;

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
```

Compile, and then load this module:

```
[vbg@vbg work test selinux]$ make
Compiling targeted test module
/usr/bin/checkmodule: loading policy configuration from tmp/test.tmp
/usr/bin/checkmodule: policy configuration loaded
/usr/bin/checkmodule: writing binary representation (version 10) to
tmp/test.mod
Creating targeted test.pp policy package
rm tmp/test.mod.fc tmp/test.mod

[root@vbg work html]# semodule -i /home/vbg/test selinux/test.pp

[root@vbg work html]# semodule -l | grep test
test    1.1.3
```

Now that our module is loaded in the memory, let us check the effect by creating a new file under the */var/www/html/* folder

```
[root@vbg work html]# touch /var/www/html/testfile.html

[root@vbg work html]# ls -Z /var/www/html/testfile.html
-rw-r--r--  root root unconfined_u:object_r:lfy_t:s0  /var/www/html/testfile.html
[root@vbg work html]#
```

Wow! It works. Now any files created in the */var/www/html/* folder will be of the type *lfy_t*. It is left as an exercise to the reader to ensure that the default contexts of all files created in */var/www/html/* should be set to *lfy_t* so that *restorecon* will set the correct type (*lfy_t*). *Hint:* modify the *test.fc* file with some regular expressions

In the next article, we will further explore type transition rules, but from the viewpoint of running processes. We will also start exploring macros and their uses. Do feel free to drop me a mail, if you face any issues while playing around with SELinux. END

By: Varad Gupta

Varad is an open source enthusiast who strongly believes in the open source collaborative model not only for technology but also for business. India's first RHCSS (Red Hat Certified Security Specialist), he has been involved in spreading open source through Keen & Able Computers Pvt Ltd, an open source systems integration company, and FOSTERing Linux, a FOSS training, education and research training centre. The author can be contacted at varad.gupta@fosteringlinux.com

A Conceptual Introduction to SCTP

Let's see what is SCTP and its advantages over TCP

The Stream Control Transmission Protocol (SCTP) is an IP transport protocol. It resides at an equivalent level with TCP and provides a reliable transport service, ensuring that data is transported across the network without error and in sequence. It is designed to transport Public Switched Telephone Network (PSTN) signalling messages over IP networks, but is capable of broader applications.

SCTP employs a session-oriented mechanism (similar to TCP) for transmission of data. A relationship is established between the endpoints of SCTP associations prior to data transfer and this relationship is maintained until all data transmission has been successfully completed.

There are many good features in SCTP, not all of which, obviously, can be covered in one article. The most important ones are:

1. Multi-stream capability
2. A broader concept for connecting endpoints (similar to, and more useful than, TCP connection), known as 'association'
3. Multi-homing support
4. Path and peer failure detection

Let's look at these in detail.

Why the name 'Stream'?

What's in a name? A lot, in this case. While TCP assumes a single-stream data transfer with byte sequence preservation, SCTP offers the multi-streaming facility. The term 'stream' in SCTP means "...a sequence of user messages that are to be delivered to the upper-layer protocol in order with respect to other messages within the same stream." This is in contrast to its use in TCP, where it refers to a sequence of bytes.

Streams allow user data to be partitioned into multiple streams, each with independently sequenced delivery. Therefore, message loss at one stream will not affect (or hold up) the delivery of other

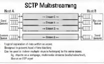
SCTP Multistreaming
Host A
Host B
Send Queues
Receive Queue
Stream 0
Stream 1
Stream 2
Stream 0
Receive Queue
Send Queues
Logical separation of data within an assoc
Designed to prevent head-of-line blocking
Can be used to deliver multiple objects belonging to the same assoc
Eg. objects on a webpage, multimedia streams (audio/video/text),
files in an FTP mget

Figure 1 Streams

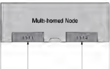

Figure 2. An example of a multi homed node

streams SCTP streams are unidirectional—see Figure 1

This reduces the delay caused by head-of-line blocking experienced in TCP since there is an independence of delivery across streams within an association. Even programmatically one finds it useful to separate logically different streams of data over the same socket, and SCTP provides this facility For those of you who exclaimed "This brings the same advantage to data transmission that threads brought to process management!"—you are right more or less.

Multi-homing

A network element can be called multi-homed if it can be addressed by and thus owns multiple IP addresses—see Figure 2

When a peer is multi-homed, one of the IP addresses is selected as the primary destination address All data will be destined to this IP by default When a failure of a primary address is detected the sender selects an alternate primary address until the link is restored or the user initiates a change in the primary address

During the process of association establishment, an SCTP endpoint can (optionally) provide a list of transport addresses (i.e multiple IP addresses in combination with an SCTP port—thus making it multi homed) through which that endpoint can be reached and from which

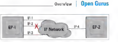

Figure 3 Example of multi-homing during network failure

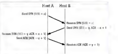
Host A
Host B
Send SYN (SEQ = x)
Receive SYN (SEQ = x)
Send SYN (SEQ = y, ACK = x + 1
Receive SYN (SEQ = y, ACK = x + 1)
Send ACK (ACK = y + 1)
Receive ACK (ACK = y + 1)

Figure 4 A three-way TCP handshake

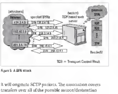

Figure 5 A SYN attack

it will originate SCTP packets. The association covers transfers over all of the possible source/destination combinations that may be generated from each endpoint's lists

This provides a better chance of the session surviving during network failures. In a conventional single homed session the failure of local network access can isolate the peer Refer to Figure 3.

SCTP provides this feature enabling sessions to be more network-fault tolerant.

Associations—more than connections

SCTP is connection oriented in nature An SCTP association is a broader concept than a TCP connection since it incorporates the multi-streaming facility and multi-homing capability.

The SCTP handshake mechanism

We are aware of the TCP 3 way handshake (see Figure 4) We are also aware of how this leads to SYN attacks wherein a server is brought to its knees, and refuses to accept any more connections (a victim of DOS attacks). Figure 5 briefly illustrates how it happens

Each time a SYN packet is received the TCP allocates/reserves a resource on its machine so a flood of SYN packets by a malicious user means that the TCP server refuses connections after the limit of its data structures is reached.

SCTP cleverly overcomes this using a four-way handshake (see Figure 6) Here SCTP plays safe and allots resources only when the connection initiating entity (client) has established its identity (and good intentions). This is how it works

1. Node A (client) creates an INIT message (an INIT *chunk* technically) and sends it to Node B (server). It also starts an INIT timer to handle timeout and kick in retransmissions.
2. If Node B is open to accept connections, it generates an INIT ACK chunk includes a unique value in it, named cookie and sends it to Node A At this point, no resources are allotted for this particular association at Node A
3. Now when Node A receives the cookie, it mirrors the same in a COOKIE ECHO chunk. An appropriate timer is started to take care of timeouts
4. Finally when Node B receives the ECHO chunk it validates the cookie sends back a COOKIE ACK chunk, and then allots resources for the association

So we see how SYN flooding leading to a DOS is overcome in SCTP design. Once the four way handshaking is done the association is said to be in an "established" state. Data transmission can legally begin now In fact, the specification even allows DATA chunks to be bundled with COOKIE and COOKIE ACK chunks.

Hello ... are you there?

Constant awareness of the peer and path status is important in SCTP This allows it to efficiently use its multi homing feature This is achieved through a Heartbeat (and Ack) mechanism A HEARTBEAT chunk is sent to any destination address that has been idle for longer than the heartbeat period The peer responds with a HEARTBEAT ACK to indicate that it is still alive Let us note that Heartbeat is sent only on 'idle' addresses On an IP that is active (i e., DATA chunks flying here and there), Heartbeat is unnecessary A doctor doesn't check the pulse of a patient to figure out whether he is dead or alive, if the patient is describing his ailment!

RFC 3286 describes the path failure detection mechanism as follows (go on read it even though it is taken from an RFC it is surprisingly interesting) "A count is maintained of the number of retransmissions to a particular destination address without successful acknowledgement When this count exceeds a configured maximum the address is declared inactive, notification is given to the application and the SCTP begins to use an alternate address for the sending of DATA chunks.

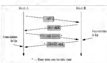

Figure 6: A four-way SCTP handshake

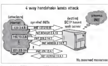

Figure 7 How SCTP overcomes DOS attacks

"Also Heartbeat chunks are sent periodically to all idle destinations (i e alternate addresses) and a counter is maintained on the number of Heartbeats sent to an inactive destination without receipt of a corresponding Heartbeat Ack. When this counter exceeds a configured maximum that destination address is also declared inactive

"Heartbeats continue to be sent to inactive destination addresses until an Ack is received at which point the address can be made active again The rate of sending Heartbeats is tied to the RTO estimation plus an additional delay parameter that allows Heartbeat traffic to be tailored according to the needs of the user application"

And as to peer failure "A count is maintained across all destination addresses on the number of retransmits or Heartbeats sent to the remote endpoint without a successful Ack. When this exceeds a configured maximum the endpoint is declared unreachable and the SCTP association is closed"

Assoc shutdown

Even in shutdown procedures, SCTP has some significant advantages over TCP For instance, SCTP implements a graceful close of an association by exchanging three messages These messages acknowledge that both endpoints will cease in their transmissions of data. This allows no room for "half open" connections (In TCP a "half-open" connection happens when one endpoint that continues to send data though the peer endpoint is no longer transmitting data.)

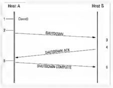

Figure 8: SCTP shutdown

There are six significant points in the shutdown-handshake scenario:

1. SCTP-user level application issues a *close()* at Node A. At this point, no new data can be sent from Node A.
2. The SCTP implementation at Node A has sent and received an acknowledgement for all queued data (before the close). At this point, the endpoint sends a SHUTDOWN chunk.
3. Node B receives the SHUTDOWN so the upper layer is notified and no new data will be accepted for transmission.
4. Node B has received acknowledgment for all its queued data so it sends a SHUTDOWN-ACK.
5. Node A destroys the association/TCB and sends back a SHUTDOWN-COMPLETE.
6. Node B receives the SHUTDOWN-COMPLETE and destroys its association/TCB.

This is a clean and graceful closure of an association. There is also an option of an 'ungraceful' shutdown using ABORT chunks. Discussions about this (and more) will be left for another day.

Acknowledgements

Randall Stewart (randall@stewart.chicago.il.us), who graciously permitted the author to use material from his wonderful informative site (www.sctp.org). Thanks, Randall!

References

- *www.sctp.org*: Informative site, contains quality material for (non-)beginners.
- RFC-3286 – An Introduction to SCTP
- RFC-4960 – Stream Control Transmission Protocol

By: Srikanta Prasanna

The author likes his beard and handmade khadi clothes. In his leisure time he works for Nokia Siemens Networks, Bangalore. He can be reached via his Gmail id: srikantap.

Tips & Tricks

Get the PID of a running process

You can check the PID of a running process/daemon by using the *pidof* command as follows:

```
grao@linux ygm3 ~> pidof program name
```

For example, to check what the PID of *sshd* is on my system, I can use the following command:

```
grao@linux ygm3:~> pidof sshd
1863
```

—Gururaj Rao, raogrr@gmail.com

System reconfiguration on the next reboot

We can change the following configurations at one go when the system boots the next time: *passwd, netconfig, timeconfig, kbdconfig, authconfig* and *ntsysv*. Use this command as the root:

```
sys-unconfig
```

...and reboot the system to reset these details.

—Vineeth T V, vnthtv@gmail.com

Using Vim as a MySQL pager

When using MySQL, you can set any pager you wish using the pager command.

```
mysql> pager less
PAGER set to less
mysql> pager vim
PAGER set to vim
```

Using Vim as your pager allows you to quickly export your query output into whatever format you want. After issuing the above command, you can see your query results displayed in Vim.

—Vishwanath Hariharan, vishi83@gmail com

Broadcast a message on all terminals

You can broadcast your message to everyone who's logged on at all terminals by using the *wall* command. Now if you want to broadcast a message to restart the system in five minutes, then open a terminal and run the following command.

```
wall system is going to restart in 5 minutes
```

The above message will be broadcast to all terminals.

—Anand Nayyar, anand_nayyar@yahoo.co.in

Simple Vim trick

Want to check the man page of a command or use a system call directly from your Vim editor? Place your cursor on the keyword and hit *Shift+k*. This takes you to the man page or the command where your cursor was, if it exists. Once you are done reading, hit *q* followed by *Enter* to get back to the Vim editor.

—Purna Chandar, purnachandar.midi@ oneconvergence.com

Translating characters into a particular case

I have struggled a lot while writing a shell script to compare some *if* conditions in order to do a string match due to case sensitivity issues. The easiest solution is to convert each character in the string by using the *tr* binary.

Converting lower case characters into upper case characters in a string:

```
echo "SomeThing" | tr [:lower:] [:upper:]
```

Converting upper case characters into lower case characters in a string:

```
echo "SomeThing" | tr [:upper:] [:lower:]
```

—M Sivakumar, it.sivu@yahoo com

Add custom start-up scripts in Ubuntu

Let's assume you want to add a custom script named *myscript* to the Ubuntu SysV *init* process. Place your script in the */etc/init.d/* directory, and make sure the file permission is 755. Now execute the following command as the root:

```
update-rc.d myscript defaults
```

Note that when you use the *defaults* option, then *update-rc.d* will make links to start the service in run-levels 2, 3, 4 and 5, and to stop the service in run levels 0, 1 and 6.

To remove the script from the init system:

```
update-rc.d myscript remove
```

This is applicable for Ubuntu and derived distros only.

—Debapriya Biswas, debapriya.biswas@gmail.com

whatis and whereis

You can find any apps path by using the *whereis* command as follows:

```
navjeet@desktop:~$ whereis pidgin
pidgin: /usr/bin/pidgin /usr/lib/pidgin /usr/share/man/man1/pidgin.1.gz
```

You can find the description of apps using the *whatis* command.

```
navjeet@desktop:~$ whatis iptables
iptables (8)   administration tool for IPv4 packet filtering and NAT
```

—Balachandar K M, balachandarlinks@gmail.com

Display the system's SMBIOS hardware components

Do you want to know the entire details of each piece of hardware on you computer? Here is the command for it (run with root permissions)

```
dmidecode -t x
```

Replace x by:

- 0 for BIOS
- 1 for system
- 2 for base board
- 3 for chassis
- 4 for processor
- 5 for memory controller
- 6 for memory module
- 7 for cache
- 8 for port connector
- 9 for system slots
- 10 for on board devices
- 11 for OEM strings
- 12 for system configuration options
- For more available options, refer man pages

—Mallikarjun, malik.v.arjun@gmail.com

Extracting a portion of a log file

Sometimes we might need to extract a small portion (100 or 200 lines) of a log file that has thousands (or lakhs) of lines. This happens when we have to send the log in a mail or while reporting a bug. Use the following command to do this:

```
sed -n '/^starttime/,/^endtime/p' file > file.part
```

For example, if you want to extract the log between 'Jul 12 10:58:26' to 'Jul 12 11:31:05' from */var/log/messages* to *messages.part*, the command would be:

```
sed -n '/^Jul 12 10:58:26/,/^Jul 12 11:31:05/p' \
/var/log/messages > messages.part
```

The "^" character is used to match the time in the beginning of a line. The format of the start time and end time depends on the log file you are working on. This works only if both the start and end times are recorded in the log file.

—Sree Pratheep, sreepratheep@gmail.com

Limit the CPU usage of a process

You can use the *cpulimit* command to limit CPU usage of any process or application. You can limit a certain running application, either by its name or by its PID. For example, to restrict the VLC media player to go beyond a 20 per cent CPU usage limit, use the command below:

```
cpulimit -e vlc -l 20
```

We can also use the PID as follows:

```
cpulimit -p 6510 -l 40
```

To find the process name or PID, use the *ps -el* command.

—Neeraj Kumar, neeraj.jnumca07@gmail.com

Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in LFY so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at *www.linuxforu.com*. The sender of each published tip will get an LFY T-shirt.

CODECHEF

The CodeChef.com Challenge—your monthly dose of coding puzzles from India's biggest coding contest, now in print!

Welcome back to the CodeChef Puzzle Challenge sponsored by *LINUX For You* and Directi. It's time to tease your brain with another interesting problem, and while you're at it, you may just win some cash as well!

Last month's puzzle

A proper divisor of a number is a divisor that is less than the number itself. Find the sum of all proper divisors for the number 1,23,748. For example, the proper divisors for the number 10 are 1, 2 and 5. So the answer would be 1+2+5 = 8.

The solution

The sum of all the proper divisors for the number 1,23,748 is 92,818.

How did we do it?

A simple solution would be to run a for loop from 1 to 1,23,747 and if the current number under consideration divides 1,23,748, then add it to the sum. If we have to find the divisor-sum for a lot of numbers, a more advanced technique like sieving could be used to pre-compute the answer.

And the winners for January are...

- Neeraj Kumar
- Debanjan Chanda
- Osric Fernandes

Puzzle for February

The Fibonacci number series is a famous one where *F[n] = F[n-1]+F[n-2]*. The first few numbers of this series are 0, 1, 1, 2, 3, 5, 8, and so on. Considering 1-based indices, what is the 999th Fibonacci number? (For example, The 7th Fibonacci number is 8.)

Looking for a little extra cash?

Then solve this month's CodeChef Challenge puzzle and e-mail your answer to *codechef@efyindia.com* by February 16, 2010. Three lucky winners will win cash of Rs 1,000 each! END

About CodeChef: CodeChef.com is India's first, non-commercial, online programming competition, featuring monthly contests in more than 35 different programming languages. CodeChef has been created by Directi as a way to continuously challenge and engage the developer community. The goal is to provide a platform for practice, competition and improvement as well as enable developers to benchmark their skills against their peers. Log on to *www.codechef.com* for the February Challenge that takes place from February 1 to 11, 2010, to win cash prizes up to Rs 55,000.

A screenshot from the Quake 4 gameplay. Screenshot [illegible] from http://www.flickr.com/people/[illegible] The image is released under Creative Commons [illegible] 2.0 Generic License.

5 Tools to Develop AAA Games

In an earlier article titled Getting Started with Game Development published in the December 2009 issue of LFY we looked at the skeletal structure of a game and had a cursory glance over some open source libraries and tools. This time around we will look at five specific open source tools that are used frequently in many AAA titles.

To begin with let us highlight the key functions within the game. Every time you run a game a piece of software that is called the game engine loads itself into the memory. The game engine is like a kernel for the game. It will manage all resources levels and your in game progression. Everything that you see, hear or do is managed by the engine. The engine itself is a collection of services that provide these capabilities.

Once the engine is loaded it uses the rendering subsystem to set your screen and start the game interface—legal screens, intros and menu systems. Using this rendering and the game logic through the game engine you begin your gaming session. Every thing that you see is taken care of by the rendering engine. So we will look at a rendering engine—OGRE 3D to be specific.

A very integral part of your game experience is also the sound design and implementation. Without good sound integration, it will be very difficult for a game to sweep you off your feet. We will look at one of the cutting-edge sound libraries out there, called Open AL.

Now since a game is an emulation of the real world, it has to have the richness and detail of the real world. There are two aspects to this: visual richness and behavioural richness. Things not only need to look real, they also have to behave real. The models and textures used in your game, plus the capabilities of your rendering engine will determine the first aspect, while the physics and AI employed in your game determine the later. We will look at Blender and how it can help you with visual richness while we will consider the Bullet Physics library to see how physics and AI can be taken care of.

Lastly, a great multi-player experience is something that adds a lot of value to your game in the current networked world. We will look at how RakNet helps in adding that bang to your multi-player experience.

Since we are looking at five tools—all in one—we will keep it simple this time and work individually with each tool at a later opportunity. Also, since we don't claim to be experts in the field and are only learning, most of the content here is derived from our researches over the Internet, rather than personal experience.

OGRE 3D

Object-Oriented Graphics Rendering Engine (or OGRE for short) is a highly modular graphics rendering engine. Its job is to interface through the available render systems to connect to the underlying graphics hardware on your computer, process the required data and send it to the display device through the graphics hardware. To put it more clearly, OGRE takes all your levels, players, along with the pretty images and sends it to your graphics card for processing. It does so by connecting to one of the two interfaces available — DirectX or OpenGL. Through one of these render systems, the data is sent to the graphics card to be processed and displayed on your screen.

OGRE is written in objected-oriented C++ with many language bindings available for developers using other languages. It is clean and easily understandable code, has an immense amount of documentation (API, wikis and forums) and a highly active developer community (forums and IRC channels), which makes it a suitable choice for all sorts of development requirements. It sure does take an experienced C++ developer with some graphics programming experience to actually go into the details and understand the inner working of the behemoth, but with a good knowledge of STL and object-oriented programming in C++, you can be well on your way to making games or other applications that require a 3D rendering engine.

OGRE is highly modular—its pluggable architecture makes adding new features a cakewalk. From custom scene graphs—that is, compositors that post process your screens to input and GUI systems—almost everything can be plugged in and used. This level of abstraction from the underlying system gives developers more time to work on the bigger picture rather than get down and dirty pushing vertices around.

OGRE is a scene graph-based engine, with all of OGRE's scene managers plugged in to the root of the engine. OGRE comes with an octree, BSP and Paging Landscape scene manager that has support for progressive LOD (level of detail)—automatically or manually created—with many others written and supported by its 10-year-old community.

This multi-platform engine runs on Linux, Mac, Windows and other platforms like the XBox360, PS3 and the iPhone, to name a few. Its flexible licence makes it an optimal choice for both Free/OSS and commercial projects. OGRE was earlier distributed under the GPL and LGPL, but from 1.7 (Cthugha) onwards, it uses the MIT License.

OGRE is purely hardware driven—it runs on your graphics hardware, dedicated or otherwise. There has always been a question of why unlike other engines (like IRRLICHT with multiple software renderers for fallback) OGRE sticks to being purely hardware driven. The simple answer is that it focuses on cutting-edge technology and supporting the latest hardware-driven features. OGRE tries to stay lean and mean, focusing on what it does best, and letting others do what they do best. It supports Vertex and Fragment programs as well as shaders written in GLSL, HLSL, Cg and assembler.

A recent addition to OGRE's feature list is the new compositing manager and its scripting language. The compositing manager gives the developers full screen post processing for effects such as blooming, blurring, noise, HDR and much more. OGRE has an animation engine with full hardware skinning support and also complete pose mixing. This enables you to blend between animation poses, giving you smoother and more realistic animation. It also comes with a highly extensible particle manager with customisable emitters and effectors.

Apart from being a wonderful graphics rendering engine, OGRE also deals with minimal windowing, memory debugging, input, GUI systems, resource loading and management, to name a few. There are many community projects that also deal with content creation as well as the export and import of pipelines for the engine.

OGRE can be seen in action in many free and open source interactive applications and games, as well as in many commercially successful projects like *Jack Keane*, *Venetica*, *Torchlight* and *Zombie Driver*, to name a few.

Open Audio Library (Open AL)

Open AL is a free, cross platform audio API. The API and coding conventions resemble Open GL, and the library has the best multi-channel 3D positional sound support. The library is currently supported by Creative Technology and Apple, with an active community backing. It is an excellent DirectSound alternative.

The library models a collection of audio sources moving in a 3D space that are heard by a single listener somewhere in that space. The basic OpenAL objects are: a listener, a source, and a buffer. There can be a large number of buffers, which contain audio data. Each buffer can be attached to one or more sources, which represent points in the 3D space that emit some sound (audio). There is always one listener object (per audio context), which represents the position where the sources are heard—rendering is done from the perspective of the listener.

The source objects point to an audio buffer, besides the velocity, position, intensity and direction of sound. The audio buffer is the audio data. The listener object contains the velocity, position and direction of the listener (mostly your in-game character). With the help of these three objects, OpenAL helps recreate natural 3D sound in the virtual world (with natural effects such as the Doppler Effect). The library is built with an extensible architecture that makes it very agile and flexible for the future as well.

The API is highly cross platform and binaries are available for the following platforms: Mac OSX, iPhone, GNU/Linux (OSS/ALSA), BSD, Solaris, IRIX, Windows, Xbox/360, AmigaOS 3.x and MorphOS. New platform support is not likely in the near future, but this in itself is an exhaustive list.

And if you are wondering what games use Open AL, the following are just some examples: ***Battlefield 2, Bioshock, Colin McRae: DiRT, Doom 3, Ghost Recon: Advanced Warfighter, Lineage 2, Quake 4, Race Driver GRID*** and ***Unreal Tournament 2004***. The Open AL is also an integral part of the id Tech 3/4/5 and Unreal Engine 2/3 engines.

The library is licensed under LGPL. To work with Open AL, look at the detailed walkthroughs on Open AL by Dev Master (***www.devmaster.net/articles.php?catID=6***).

Blender

Once an in-house production software and later freed by its creators, Blender is the animation/CG industry's new kid on the block. Blender is a multi-platform 3D graphics application used for modelling, texturing, rigging, animating, etc. Its source code was opened and it is now released under the GNU GPL. It is also a very powerful tool when it comes to UV unwrapping, 2D/3D painting, skinning, sculpting, water creation, physics particle simulations... and no, the list doesn't quite end yet, but goes on to include post production tools as well, like a non-linear editing suite with a compositing tool to boot. Blender can also be used to create interactive 3D applications and games. It is available on multiple platforms including Linux, Windows, and Mac OSX. There are official and unofficial ports to other popular operating systems as well.

Blender has a light footprint and an intuitive interface, and covers almost all features of its commercial counterparts, even doing a better job at some places. Blender goes a step further to be a single resource for all your 3D production and post-production needs. It treads into real-time rendering and games with its integrated game engine. Not stopping there, Blender provides game logic, real-time physics and a workflow that connects your asset creation pipeline to your post-production and extends it to the realm of game development.

If you have worked on games or animation projects, you will have seen that both share a very similar development/production pipeline, only dividing towards the end, where one is rendered offline and the other in real time. So it is an obvious choice to extend a well-built 3D production and animation suite to game development. Most of Blender is written in C/C++ and Python. It can be extended using plug-ins and can be customised using Python scripts as well.

Blender sports many of the features that are on par with 3D CGI production powerhouses, like Maya and Softimage. It supports a variety of geometric primitives including polygon meshes, Bezier curves, NURBS surfaces, metaballs, fast subdivision surface modelling, digital sculpting, and outline fonts. It supports key-framed animation tools including inverse kinematics, armature (skeletal), hook, curve and lattice-based deformations, shape keys (morphing), non-linear animation, constraints, vertex weighting, soft body dynamics including mesh collision detection, LBM fluid dynamics, bullet rigid body dynamics, particle-based hair, and a particle system with collision detection. Blender renders geometry in all its glory – lit, textured and composited with the integrated YafRay ray tracer.

It has an intuitive interface, but it still can be a bit of a steep learning curve for even those who have experience in other 3D graphics applications like Maya or 3D Max, apart from those who have no previous experience. It will take almost the same amount of time to master Blender as it will to get comfortable with Maya and Max.

The Blender foundation tries to stay focused on creating cutting-edge tools and adding new functionality in Blender, while trying to make the learning process easy by conducting workshops, releasing DVDs, updating the wiki, tutorials and other documentation. The ever-growing community of Blender artists fuels the fire that burns in the hearts and minds of everyone involved with Blender.

Bullet Physics

Bullet Physics is a professional collision detection, rigid body and soft body dynamics library. The library is free for commercial use under the ZLib License. The library is under active development and its roadmap includes support for CUDA, OpenCL, and the fracture of rigid bones.

The library is a cross-platform library and supports Playstation 3, Xbox 360, Wii, PCs with the Linux and Mac OSX, and the iPhone. It has a very advanced feature set.

- Discrete and continuous collision detection including ray and convex sweet test. Support for all basic primitive shapes and concave/convex meshes.
- Fast and stable rigid body dynamics, vehicle dynamics, character control and slides, hinge physics, plus twist constraint for ragdoll physics.
- Soft body dynamics for cloth, hair/fur, rope objects and deformable volumes, supporting two-way interaction with rigid bodies plus binary constraints.
- Integration with major development tools such as Maya, Blender and COLLADA.
- Completely written in C++ and the source has been open sourced under the Zlib License.

As Bullet is fairly unknown in the general gamer circles (unlike Blender, OGRE and Open AL), you may be wondering which projects have used it. How about *Grand Theft Auto 4*, Sony's *Free Realms* and *2012* (last year's Hollywood blockbuster). Bullet Physics is also included within the RAGE Engine of Rockstar Games.

RakNet

RakNet is a cross-platform, open source, object-oriented networking library for game programmers, written in C++ by RakkarSoft. A lower-level framework powers the higher-level features such as game object replication, voice chat, and patching binaries.

It is a DirectPlay alternative touted as being better than the DirectX component due to its simpler integration. The library uses UDP for all network connections and is written primarily for game development. The library is available in three different licensing models: GPL, shareware and commercial licences. RakNet also offers free licences for indie developers.

The major features of the library include:

- Object replication: Automatically creates, destroys, serialises and transmits game objects
- Player lobby: Database-powered lobby with support for friends, rooms, quick matches, and ranking
- Secure connections: Support for SHA1, AES128, SYN cookies and RSA
- A robust communication layer with automatic flow control, message ordering on multiple channels, message coalescence, and splitting and reassembly of packets
- An autopatcher to automatically patch the software at the gamer's end with the latest updates
- Support for the various C/C++ remote procedure calls
- Support for voice communication, including bindings for Port Audio, FMOD, and DirectSound, besides NAT Punchthrough for voice support over P2P applications

As it's also used as a commercial tool, most applications of RakNet are protected under NDAs. So a finite list of games using RakNet is not freely available. However, RakNet is being used in a number of games by *Epic Games China*, *Codemasters*, and the *Unity Engine*. It was also used in the games *Crashday* and *Hyperball Racing*.

In a nutshell

Yes, yes, the tools look totally awesome, we know... and they sound jam packed with features and efficiency. Yes! You will be surprised with how frequently they are used in major titles. We suggest you check out the *Further Reading* section and try out these tools for yourself.

Step up and reach out—all the awesomeness is within reach. Sing the song!

Further Reading

- OGRE 3D Website: http://www.ogre3d.org/
- OGRE 3D Features: http://www.ogre3d.org/about/features
- The Story of OpenAL: http://www.linuxjournal.com/article/4400
- About Open AL - http://connect.creativelabs.com/openal/default.aspx
- List of Games using Open AL - http://connect.creativelabs.com/openal/OpenAL%20Wiki/Games.aspx
- Applications of Open AL (Wikipedia) - http://en.wikipedia.org/wiki/OpenAL#Applications
- Blender Website: http://www.blender.org/
- Blender 3D Features - http://www.blender.org/features-gallery/features/
- Community of Blender Artists: http://blenderartists.org/
- Bullet Physics Features List: http://bulletphysics.com/Bullet/BulletPhysics2.74_2.htm#_Toc223854815
- Bullet Physics on Google Code: http://code.google.com/p/bullet/
- The Bullet Physics Issue List - http://code.google.com/p/bullet/issues/list
- Working with Bullet Physics: http://bulletphysics.com/Bullet/BulletPhysics2.74_2.htm
- RakNet Features List: http://www.jenkinssoftware.com/raknet/index.html
- Visual Guide to RakNet Features: http://www.rakkarsoft.com/raknet/manual/RakNetMajorFeatures.html
- RakNet Licensee List: http://www.jenkinssoftware.com/licensees.html

By: Yadu Rajiv, Kinshuk Sunil

The authors are independent game developers and founders of the Indie GameDev India Community (indiegamedev.in). Their desk job takes them to OSSCube, where Yadu is a UX Engineer and Kinshuk takes care of the community initiatives. Reach them at yadu@osscube.com and kinshuk@osscube.com

Getting Started with GNU Octave

Part 3

In the first two parts of this series, we learnt how to use Octave for linear algebraic problems. We also took a quick look at solving simple optimisation problems. Now, we will explore the other features that facilitate numerical and scientific computing in Octave.

Let's delve into Octave's features that relate to statistics, polynomial manipulations and more.

Statistics

Octave has various features for the analysis and interpretation of multi-dimensional numerical data. As is true for most other cases, the data that you want to analyse is usually specified as a matrix, where a row is one observation, and a column indicates the value of the individual components of an observation.

Descriptive statistics

Descriptive statistics give a quantitative measure of the main features of a data set, such as its mean, median, variance, kurtosis and others. Let us now look at the corresponding functions in Octave:

- **mean.** The function to calculate the mean of a data set is *mean(x,dim,opt)*. If x is a vector, then it returns the mean of the elements of x, i.e.,

$$\frac{1}{N}\sum_{i=1}^{N} x_i$$

If x is a matrix, the default behaviour of Octave is to return a row vector in which each element is the mean of the corresponding column of x. For example:

```
octave-3.2.2:53> x = [1 4 3]
x =
1   4   3

octave-3.2.2:54> mean(x)
ans = 2.6667
octave-3.2.2:55> x = [1 2 3; 4 5 6]
x =

1   2   3
4   5   6

octave-3.2.2:56> mean(x)
ans =

2.5000   3.5000   4.5000
```

The optional argument *opt* may be used to specify the type of mean. The default is *arithmetic mean*, indicated by an 'a'. It is also possible to calculate the *geometric mean*, specified by a 'g', or the *harmonic mean* specified by an 'h'.

The optional argument *dim* is used to specify the dimension of the matrix along which the mean is to be calculated. By default, the dimension is set to 1. You may also specify the dimension to be a maximum of 2 for a two-dimensional matrix. When you specify *dim = 2*, the mean is calculated *along the rows*. So the mean is a column vector. For example:

```
octave-3.2.2:88 > x = [1 8 3;4 9 6]
x =

   1   8   3
   4   9   6

octave-3.2.2:89 > mean(x, 2)
ans =

   4.0000
   6.3333
```

If you have a higher-dimensional matrix, you may specify dim accordingly.

- **median:** The syntax and semantics of the *median* function are similar to the *mean* function: *median(x,dim)*, where *dim* has a similar interpretation as above.
- **kurtosis:** *kurtosis* is the degree of how much a probability distribution peaks. In Octave, it can be found by using the kurtosis function. The syntax is similar to the *median* function introduced above.
- **statistics:** If you want most of the descriptive *statistics* on your data, use the statistics function. If *x* is a matrix, it returns a matrix with the minimum, first quartile, median, third quartile, maximum, mean, standard deviation, extent of being skewed and kurtosis of the columns of *x* as its rows. If *x* is a vector, it is treated as a column vector.

Besides the above, you also have functions for calculating the *covariance*, *correlation*, *extent of being skewed* and *moment* of a set of multi-dimensional data.

Basic statistical functions

- *Mahalanobis distance* is the distance measure that is useful for identifying the similarity of an unknown sample of data to a known one. In Octave, you can use the function *mahalanobis(x,y)* to find the Mahalanobis D-square distance between the multivariate samples *x* and *y*, which must have the same number of components (columns), but may have a different number of observations (rows).
- *Miscellaneous functions:* The *perms* function may be used to generate all permutations of a matrix or a vector, one row at a time. *values* is a useful function that returns the different values in a row or a column vector in ascending order—i.e., it does not return the duplicates. Please note that values do not work on a matrix. There are also other functions of interest to a student of statistics, such as *logit* and *probit*.

Other features: Besides the rather basic features described earlier, there is extensive support for performing several different statistical tests such as equal mean values, equal median values and others. A large number of probability distributions is available using which you can generate random numbers and matrices.

Please look up the Octave reference manual under the statistics section for an extensive index of the functions available for the various statistical analyses of your numerical data.

Polynomial manipulations

Let us now learn to work with univariate polynomials in Octave. A polynomial is represented by a vector containing its coefficients. For example, a polynomial such as $p(x) = 2x^3 + 3x^2 + 4x + 5$ will be represented in Octave by the vector c = [2 3 4 5]. The *polyout* command prints the formatted polynomial from a coefficient vector. Let's try this out.

```
octave-3.2.2:92 > c = [2 3 4 5 ];

octave-3.2.2:93 > polyout(c, 's')
2*s^3 + 3*s^2 + 4*s^1 + 5

octave-3.2.2:94 > polyout(c, 'x')
2*x^3 + 3*x^2 + 4*x^1 + 5
```

As you can see, we have first defined the coefficient vector *c*, and then used the *polyout* function to print the formatted polynomial. The second argument to *polyout* is the variable name in the univariate polynomial.

Roots and evaluation

The *roots* function returns the roots of the polynomial, given its coefficient vector *c*:

```
octave-3.2.2:100 > r = roots(c)
r =

  -1.37119 + 0.00000i
  -0.06443 + 1.34878i
  -0.06443 - 1.34878i
```

Now let's look at how good the roots are. Let us evaluate the polynomial at the obtained roots.

```
octave-3.2.2:101 > polyval(c, r(1)(1))
ans =  2.6645e-15
octave-3.2.2:102 > polyval(c, r(2)(1))
ans =  6.2172e-15 + 3.6915e-15i
octave-3.2.2:103 > polyval(c, r(3)(1))
ans =  6.2172e-15 - 3.6915e-15i
```

Okay, so it looks like the roots are pretty good, as the value of the polynomial is tending to 0 at those points. The *polyval* takes the coefficient vector as the first argument and the point at which the value is to be evaluated as the second argument. Instead of passing the roots one by one, as we have done, we could also send the whole matrix to *polyval*, as follows:

```
octave-3.2.2:104 > polyval(c, r)
ans =

   2.6645e-15 + 0.0000e+00i
   8.8842e-16 + 3.8527e-15i
   6.2172e-15 - 3.6915e-15i
```

As a general rule, when a matrix or a vector is passed as the second

You **said it...**

Request for articles

I have been a regular reader of *LFY* for the last three months. I find the content and presentation really good (the DVD content included). I have just started out in the Linux world and follow the articles for admins regularly. In the December 2009 issue, you had mentioned that Nagios was covered in detail in December 2007. I would like to request you to send me a PDF of that article. Also, I would like to know whether Openfiler and Zamanda have been covered earlier.

—*Debojyoti Gope, gope.debojyoti@gmail.com*

ED: *Glad to hear that you find the content useful. By the time you read this issue, we hope you'd have already read the Nagios article sent to you. As for the other two topics you wanted to know about, well, although we've not covered either, we'll definitely try to do so in the coming months.*

I'm a regular reader of your publication. I had mailed you earlier about including free computer e-books on programming languages, networking, operating systems and security, if possible, on a DVD, along with the various operating systems and FOSS products you provide—since those are any techie's real requirements.

I sincerely believe you guys should include a valuable column by a reputed software professional, about how to make money and subsequently a good living, as an open source developer. I have always wondered how one can make money through open source. If one works hard to build software and then goes on to open source its code, where would the remuneration come from to support a family? I don't understand how open source developers earn their bread and butter.

—*Bonny Davis, Haryana, davisbonny@gmail.com*

ED: *Thanks for a detailed mail. I guess you've hit the nail on the head. A column/article addressing your concerns will not only help you, but could also guide many students/software professionals to consider FOSS seriously. We'll definitely work this out in the coming months.*

Replies to an earlier post, titled: "We don't deserve the pride of being open"

Here is my reply to Vaibhav Kaushal's letter published in the January 2010 issue.

I find his experiences highly one-sided. There are many occasions in my professional career when I've faced precisely the opposite kind of situation. Take my graphical software SWISHMAX, which works fine on Windows. But on numerous occasions it halts because of viruses, et al. For the last two years, I have booted into Linux and worked on SWISHMAX using Wine, without fear of viruses, slow downs, hang ups, etc.

As for antivirus software for Windows—let's just say I'm thoroughly sick of it! Internet Explorer 5, 5.5, 6 and 7 are all slow and bulky. Firefox on Linux does everything that IE does on Windows. And it is fast and efficient on all counts.

The other day one of my assistants made a DOCX and XLSX file on Office 2007. His presentation failed to open on an earlier MS Office package. The lecture session almost fell flat, right in the presence of a waiting audience; thankfully I could open those two files on OpenOffice.org 3.1.1 on my Mandriva 2009 laptop. It shows that an OOo can do what earlier versions of MS Office failed to notice.

All my websites have run without any problems since the last five years on LAMP platforms. This includes one LAMP stack that's run since 2005 on FC-3. Had it been on IIS, I would not have had the time to write this letter.

The MS Frontpage is a good package for learning Web page designing. But for a professional Web page, you have to throw it into a dustbin and use Bluefish or Texmaker.

PHP is a lot faster and sleeker then the .NET or VB Script—don't take my word for it; check it yourself.

Even my organisation, which is a happy user of Windows since many years, relies on Linux to run its e-mail server on Lotus Domino.

It's true that there are lots of difficulties yet to be overcome, but we users should embrace these difficulties as opportunities for further development and move ahead.

Every difficulty is like a dagger. There are two ways to handle this situation: catch the dagger by the blade and get hurt, or catch it by its handle and use it as a weapon to move ahead. I would like to be in the second situation.

—*S. Bera, Mumbai, berasomnath@gmail.com*

FOSS is not just about free software, just like *Medecins Sans Frontieres* (Doctors without Borders) is not just free medicare. Both are among the noblest practically proven concepts of our time. Those who value software freedom are indebted to the selfless efforts of the FOSS community. The software-hardware racketeering of the proprietary sector is only too well known. FOSS deserves active government and private sector encouragement and support—which, we all know, is generally lacking.

In the face of all this, FOSS has been expanding and accelerating its march forward. It charts its own course as to what has to be developed and *gifted* to the world in the genuine interest of users. It is the proprietary sector that is continually watching FOSS, and not the other way around.

The proprietary sector initially chose to ignore FOSS as a *non-entity*. Today it is too late for them to try their argument to *polyval*, the polynomial is evaluated at each of those points. Let us now redefine *r* as a square matrix and see the results:

```
octave-3.2.2:108> r = [1 2; 3 4];
octave-3.2.2:109> polyval(c, r)
ans =

   14   41
   98  187
```

Polynomial curve fitting

The idea of polynomial curve fitting is that, given a set of input and output vectors, the best curve passing through the given points needs to be fitted. This is usually done by minimising the least squares distance between given points and the curve. For more background on this subject, please see the Wolfram Mathworld article given under *Resources*.

The *polyfit* function is used to find the best-fit curve for a set of input and output vectors. I shall illustrate this with an example from the MATLAB documentation of the same function (apologies for the digression, but this also illustrates the compatibility between MATLAB and Octave, a topic that we shall discuss in more detail in a later article):

```
octave-3.2.2:113> x = [1 2 3 4 5];
y = [5.5 43.1 128 290.7 498.4];

octave-3.2.2:114> p = polyfit(x, y, 3)
p =

  -0.19167   31.58214  -60.32619   35.34000

octave-3.2.2:115> x2 = 1:.1:5;
octave-3.2.2:116> y2 = polyval(p, x2);
octave-3.2.2:117> plot(x, y, 'o', x2, y2);
octave-3.2.2:118> grid on
```

You should now see a plot similar to the one shown in Figure 1. Okay, so what we have done here is obtained the coefficients *p* of a best-fit curve between *x* and *y* with *polyfit*. Next, we used the obtained coefficients to get the output values *y2* for the same x values over a finer range *x2*, and then plotted the results along with our earlier *x,y* (indicated by the circles). As you can see from the plot, they are almost on the obtained curve.

By default, *polyfit* just returns the coefficients of the polynomial. However, if you also give a second return argument, like *[p, s] = polyfit(x, y, 3)*, then it returns a structure of values containing the following fields:

- **R:** The Cholesky factor of the Vandermonde matrix used to compute the polynomial coefficients.
- **X:** The Vandermonde matrix used to compute the polynomial coefficients.
- **df:** The degrees of freedom.
- **normr:** The norm of the residuals.
- **yf:** The values of the polynomial for each value of *x*

Other features: There are also various other functions to work with polynomials, such as *polyderiv* and *polyinteg* that return the coefficients of the derivative and the integral polynomial of the function, respectively.

Miscellaneous features

Here are some other facilities available in Octave that you can explore on your own.

Plotting: Briefly used earlier, Octave has got good facilities for plotting data in two and three dimensions. You can also interact with *gnuplot* from within Octave, which is a very good feature to have.

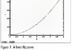

Figure 1: A best fit curve

Sets: Octave has some support for working with sets. A *set* is created with the *create_set* function. Set operations—such as *union*, *intersection*, *ismember*—are also supported in Octave.

Control system toolbox: The control system toolbox looks very well developed, though I have not explored or studied it in depth.

Image and audio processing: Octave has got some support for loading, display and some basic operations for images. There is also some support for audio processing.

Looking ahead

In the next article in this series, we will learn to write Octave scripts and take a step towards using Octave as a programming language rather than as a command interpreter, as we have been doing so far. See you next time!

Resources

- GNU Octave: *www.gnu.org/software/octave/*
- GNU Octave Documentation: *www.gnu.org/software/octave/docs.html*
- GNU Octave Function reference: *octave.sourceforge.net/doc/index.html*
- Least Squares Fitting: *mathworld.wolfram.com/LeastSquaresFitting.html*

Acknowledgements

Considering the heavy usage of Math formulas, LaTeX was an automatic choice to write this article and thanks to the *The Not So Short Introduction to LaTex* (http://ctan.tug.org/tex-archive/info/lshort/english/lshort.pdf) for a nice treatment of how to write Math well in LaTex.

By: Amit Saha

The author blogs at http://amitksaha.wordpress.com and he welcomes queries and suggestions at amitsaha.in@gmail.com

SPECIAL OFFER

SUBSCRIBE

Lifeline of Cricket
Cricket today

Get exiting Gifts

+

TIMEX Rs 1500/-

Allen Solly Rs 749/-

Rs 395/-

FREE

It's a profit game

Yes, I would like to subscribe Diamond Cricket Today english for the following term

☐ 3 Year (36 Issues). I pay Rs. 1800/- get Timex watch worth Rs. 1500/+ an IPL book worth Rs. 395/- FREE!

☐ 2 Year (24 Issues). I pay Rs. 1200/- get Allen Solly tie worth Rs. 749/+ an IPL book work Rs. 395/- FREE!

Name Mr/Ms .... Date of birth □□/□□/□□
Address ....
.... PIN□□□□□□
Occupation ....
Ph .... E-mail ....
Please find enclosed ( ) Cheque/DD No .... Dated ....
for Rs 1800 □ Rs 1200 □ favouring **Diamond Magazines Pvt. Ltd.**
Please charge it to my □ Visa □ Master Card No □□□□□□□□□□□□□
Card expiry date□□/□□ Verification Code □□□ (This code is the last three digits of the number printed at the back side of your credit card)
Date .... Signature ....
Please fill in this order form and mail it to

Cricket today

Diamond Magazines Pvt. Ltd., K-30 Okhla Industrial Area, Phase-II,
New Delhi-20 Ph 41611861 Fax. 41611866 Email: crickettoday@dpb.in

Please note: The gift may vary from the one shown above.

Hudson

An Extensible Continuous Integration Engine

This article discusses the usefulness of the Hudson CI engine, followed by how to set it up and configure jobs to perform various tasks automatically or when required.

If you believe in agile software development practices, you already realise the importance of continuous integration. The agile method is to facilitate the delivery of bug-free software to users. The following practices define the truly continuous integration system:

- Members integrate source code frequently
- Each integration is verified by an automatic build
- The code is tested in a mock production environment

The effects of using the continuous integration system are:

- Integration issues are discovered early
- The status of the latest build is easily known
- The current build is available instantly for demonstration or release.

Figure 1 illustrates the architecture of a continuous integration system.

Introducing Hudson

The Hudson system [*http://hudson-ci.org*] is designed precisely to meet continuous integration needs. It is useful to build and test software projects continuously. It provides an easy-to-use interface that makes it easier for developers to integrate changes to a software project. An automated and continuous build increases the productivity of a team.

It is also useful to monitor the externally-executed builds and notify interested parties when something is wrong. Developers can save the changes in the version control system (VCS). The Hudson job(s) monitor the VCS continuously and act accordingly.

Figure 2 shows a screenshot from the Apache build farm that hints how the Hudson home page would look after configuring multiple jobs.

The following are some of the important features of Hudson:

- **Easy installation:** There're no additional installation instructions nor any need to set up a database separately. Just download the latest stable version of the software (*hudson.war*) from *http://hudson-ci.org* and install it as follows:

  ```
  java -jar hudson.war
  ```

- **Java and open source:** Hudson is fully written in Java and hence it's platform independent. It's released under the MIT licence.
- **Easy configuration:** You can

configure it completely from a Web based GUI using your browser. There is no need to tweak or configure config files manually.

- **Distributed builds:** Hudson can distribute build/test loads to multiple computers—it uses hardware resources effectively.
- **JUnit/TestNG test reporting:** The Hudson system can parse the JUnit/TestNG test output generated by build/test scripts and generate the intuitive report.
- **Plug-ins:** You can extend its functionality by using various plug-ins. You can write your own plug-ins to support the tools and processes you use, or visit the following link to view all available Hudson plug-ins: *wiki.hudson-ci.org/display/HUDSON/Plugins*

Jobs

Builds in Hudson are configured as 'jobs'. They may involve one or many tasks, including code compilation, executing unit tests, publishing build artefacts, updating issue trackers, identifying code violations, generating documentation, etc.

The build can be configured as a series of shell, Perl and Python commands. The best method is to write a script using any scripting language that performs all activities and execute the script from a Hudson job. Depending on the plug-ins installed, you may see many options in the 'job configuration' screen. The job may belong to any of following categories:

- *Scheduled* - The job can be scheduled to run at a specific date/time—*crontab*-like scheduling is possible.
- *Monitor VCS* - You can configure a job to run immediately after a commit in the VCS. The Hudson system can poll the VCS periodically, detect the changes, if any, and initiate the build process.
- *On demand* - A build can be triggered on a demand-basis.

Figure 3 demonstrates the Job Dashboard. As you can see, it displays details about the job—including test trend, disk usage report and violations (if any).

In short, you can configure a job to update the VCS working copy, execute the build, generate the installation packages, run tests on different platforms, generate documentation and code quality reports. The final result is published in the user interface.

Plug-ins

You can manage the various plug-ins you require from within the user interface. All users have the ability to keep track of installed and new plug-ins. In fact, you can track the updates to Hudson and the plug-ins from the same screen. Users are notified of updates automatically.

Here are some of the tasks for which you might want to consider the use of plug-ins.

- Leader board
- Report code violation
- Webtest plug-in
- Publish documents
- Integration with FTP/SCP servers
- Integration with issue trackers

You can configure Hudson to run multiple slave instances in parallel, on different platforms. We will discuss the Hudson master-slave set-up in the following section.

Master-slave set-up

You can schedule your tasks to run on the same machine (master) or on a different machine (slave). A master's job is to manage one or more slaves. However, remember that although a master serves HTTP requests, it can still build projects on its own.

Slaves are computers that are set up to build projects triggered by the master. A separate program called a 'slave agent' runs on a slave computer. Let us now understand how to set up Hudson to execute distributed builds using a typical master-slave set-up. Note that you can configure a single

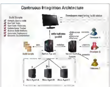

Figure 1: A continuous integration architecture

computer to execute multiple slave agents.

The master-slave set-up is very helpful when you need to execute the job on different types of operating systems. Consider a situation in which the application is expected to run on Linux, Solaris and Windows. To address this need, you can install Hudson on any machine—with Linux, for instance—and add two slaves with Solaris and Windows. The user can add three different jobs—one running on the master and others running on slave machines.

When slaves are registered to a master, the master starts distributing loads to slaves. The delegation depends on the specific job. You can either configure a job to execute on the master, or tie it to a specific slave. On the other hand, you can configure jobs to freely roam between slaves—i.e., the job will execute on the slave that's free of any load. For the user, this set-up is transparent—you can view results for all jobs using the master, irrespective of the slave that executed the jobs.

You can launch a slave from a master using any of the following methods (Figure 4):

1. Launch slave agents on UNIX/Linux machines using SSH
2. Launch slave agents via JNLP
3. Launch slave agents by executing a command on the master
4. Let Hudson control Windows slaves as a Windows service

Launching slaves on UNIX via SSH

Hudson has a built-in SSH client implementation that can be used to talk to remote sshd daemons to start a slave agent. This is the most convenient and preferred method for UNIX slaves, which normally have SSH servers running out of the box.

Click *Manage Hudson→Manage Nodes→New Node*. In this set-up, the connection information is supplied—including the SSH host name, user name and authentication credentials. You'll need to authenticate using the password of the remote (slave) UNIX system, or SSH keys. If it is configured to use SSH keys, the SSH public key should be copied to the *~/.ssh/authorized_keys* file. Hudson will do the rest of the work by itself, including copying the binary needed for a slave agent, and starting/stopping slaves.

Figure 2: The Hudson home page after configuring multiple jobs

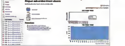

Figure 3: The job dashboard

Figure 4: List of launch methods for slaves

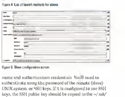

Figure 5: Slave configuration screen

In short, when you launch a slave using this method, a new Java process is started in the slave automatically. Isn't it the most convenient set-up on UNIX systems?

> **Tips:** Depending on the project and the availability of hardware resources, someone's desktop can be used as one of the slaves, without affecting his or her day-to-day activities. This will avoid the requirement for dedicated slaves.

Launching the slave agent via Java Web Start

Another way of launching a slave is to start a slave agent through Java Web Start (JNLP). In this approach, the user will log in to the slave node, open a browser and open the slave page by pointing to the URL of the master system. It may look like the following URL: *http://masterserver:PORT/hudson/jnlpJars/slave.jar*

You are then presented with the JNLP launch icon. If you click on the icon, the Java Web Start kicks in and it launches a slave agent on your computer.

This mode is convenient when the master cannot initiate a connection to slaves, such as when it runs outside a firewall while the rest of the slaves are within the firewall. The disadvantage is, if the machine with a slave agent goes down, the master has no way of re-launching it on its own.

Launching a slave agent 'headlessly'

This launch mode uses a mechanism very similar to Java Web Start, except that it can run without using a GUI—making it convenient to execute as a daemon on UNIX systems. To do this, you should configure this slave to be a JNLP slave by downloading *slave.jar* and then from the slave, run the following command:

```
java -jar slave.jar -jnlpUrl \
http://masterserver:PORT/computer/slave
name/slave-agent.jnlp
```

The *slave.jar* file is downloaded from the above-mentioned URL. Make sure to replace *slave name* with the name of the slave set-up in the master.

> **Tips:** By default, Hudson runs on port 8080. It can be installed and managed without the need for super user privileges.

Launching a slave agent using your own scripts

If the above modes are not flexible, you can even write your own script to launch a slave agent. The script is placed in the master computer and Hudson runs this script whenever it should connect to the slave. The script may use a remote login program like SSH or RSH to establish a connection between the master and slave.

The script would execute the slave agent program like *java -jar slave.jar*. The STDIN and STDOUT for the script should be connected to the master. For example, the script that executes *ssh myslave java -jar ~/bin/slave.jar* from the master Web interface would satisfy this need. For precisely this reason, running this script manually from the command line does no good.

Launching a slave agent using this mode requires an additional set up on the slave. The benefit is that when the connection goes bad, you can use the Hudson Web interface to re-establish the connection.

> **Note:** You should update *slave.jar* every time the Hudson system is upgraded. Though, in practice, *slave.jar* is not changed frequently.

Launching slave agents on Windows machines as a service is out of the scope of this article.

Hopefully the above discussion has intrigued you enough to make you configure a simple Hudson job that performs certain tasks that are specific to a build and a test. Once you start using it, you will definitely be interested in its scope to execute distributed builds.

By: Bhuvaneswaran A and Nithya Pattabiraman

Bhuvaneswaran is part of the Production Engineering Group at CollabNet and contributes to Ubuntu and Subversion projects during his spare time. You can find more about him at http://livecipher.com. Nithya is part of the Quality Assurance team at CollabNet and you can reach her at nithya@collab.net

Python in Research, Part 6

Discrete Event Simulation

Many problems in real life are very complex. One can construct a mathematical model to understand them. Yet in some cases solving the models perfectly is virtually impossible. This is when numerical experiments or simulation can be very useful SciPy includes a large number of statistical distributions.

For discrete event simulation SimPy is a very interesting and useful application. It's not yet integrated with SciPy/NumPy but, hopefully it soon might be.

In this article lets simulate queuing delays in a bank, building on top of the excellent tutorial included with the SimPy documentation (see the *References* section). The exercise can be extended by using statistical distributions from the SciPy package along with its simulation framework.

Statistical distributions

You can explore the Erlang distribution [*http://en.wikipedia.org/wiki/Erlang_distribution*]—a good option for modelling the customers' arrival time and the time spent servicing each one. It is still widely used in the telecommunications industry Try the following code:

```
import matplotlib.pyplot as plt
import scipy as np
from scipy.stats import erlang
# 3 variations with mean = 2
er1=erlang(1,scale=2)
print 'mean, variance', er1.stats()
er2=erlang(2,scale=1)
print 'mean, variance', er2.stats()
er3=erlang(10,scale=.2)
print 'mean, variance', er3.stats()
# plot the probability distribution fn
x = np.linspace(0,5)
p = plt.plot(x, er1.pdf(x))
p=plt.plot(x, er2.pdf(x))
p=plt.plot(x, er3.pdf(x))
plt.show()
```

The plot in Figure 1 shows what the probability distribution functions look like for the Erlang distribution with the parameters chosen. The function is like an exponential distribution when the shape parameter is 1 and seems to be a good option to understand the pace of customer arrivals. It looks a lot like a normal distribution for larger values of the shape parameter. However the distribution with a shape parameter of 2 looks like what you may have often experienced—no one leaves quickly and some seem to be at the counter for a long time. All three distributions have the same mean but their variances differ. The output of this program will be as shown below:

```
mean, variance (array(2.0), array(4.0))
mean, variance (array(2.0), array(2.0))
mean, variance (array(2.0), array(0.4))
```

A SimPy model

Your model consists of a source that generates customers randomly. Each customer waits for a counter (resource) to become available. The customer then spends a random amount of time to get his work done before leaving. We'll include the necessary packages and write the code for generating customers:

```
import SimPy.Simulation as simpy
from scipy.stats import erlang
class Source(simpy.Process):
  """ Generate customers randomly"""
  def generate(self, number, resource, pdf, pdf_counter):
    for i in range(number):
      c = Customer(name = "Customer%02d"%(i,))
      simpy.activate(c, c.visit(pdf=pdf_counter, counter=resource))
      # wait for the next customer
      t = pdf.rvs()
      yield simpy.hold, self, t
```

The class *Source* contains one method—*generate*. It takes the following as parameters: the number of customers to be generated, counter resources and probability distribution functions for the arrival of customers and the service time at the counter. You create a customer and activate the customer by passing a method name (*visit*, in your case) as a parameter. The simulation package will call the visit method, which will need to be defined in the *Customer* class. Then, you find the random time at which the next customer is expected and hold the process till then.

You may now define the *Customer* class as follows:

```
class Customer(simpy.Process):
  def visit(self, pdf, counter):
    time_in = simpy.now()
    yield simpy.request, self, counter
    wait_time = simpy.now() - time_in
    tib = pdf.rvs()
    yield simpy.hold, self, tib
    yield simpy.release, self, counter
    time_in_bank = simpy.now() - time_in
    process(self.name, time_in, wait_time, time_in_bank)
```

The *Customer* class contains the method *visit*, which encompasses the logic for the interaction of the customer with the bank. You track the time the customer comes in and as he waits for a counter to become available, you compute the wait time. Then, you find the random time for servicing this customer and hold the process till it gets over. After this, you release the counter. Finally you'll compute the time spent by the customer in the bank and process the times.

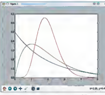

Figure 1: Erlang distribution probability functions with mean 2. Shape = 1(Blue) Shape = 2(Green) Shape = 10(Red)

```
def main(maxTime, maxNumber, arrival_interval, mean_serve_time, counters=1):
  print("%d counters with mean arrival and service times %4.1f %4.1f minutes"
      %(counters, arrival_interval, mean_serve_time))
  k = simpy.Resource(capacity=counters, name="Counter", unitName="Clerk")
  simpy.initialize()
  s = Source('Source')
  simpy.activate(s, s.generate(number=maxNumber,
                pdf=erlang(1, scale=arrival_interval),
                pdf_counter=erlang(2, scale = mean_serve_time/2.0),
                resource=k), at=0.0)
  simpy.simulate(until=maxTime)
```

Your main program starts the simulation by creating the necessary resources and activating the source process for generating customers using the *generate* method defined in the *Source* class. You use the Erlang distributions with shape factor 1 and 2 for arrivals and serving time.

```
def process(name, time_in, wait_time, time_in_bank):
  print("%s %6.2f %6.2f %6.2f"%(name, time_in, wait_time, time_in_bank))

main(100, 5, 1.5, 2.0, 1)
```

For processing, just print the results and run the simulation with a customer arriving every 1.5 minutes and needing an average of 2 minutes to be served. You would expect the queue to build up pretty quickly as you can see from the results:

```
1 counters with mean arrival and service times  1.5  2.0 minutes
Customer00   0.00   0.00   1.80
Customer01   4.36   4.73   3.21
```

```
Customer02  0.41  3.17  5.52
Customer03  0.44  5.48  7.86
Customer04  1.40  6.90  0.00
```

Here is another run and the counter is usually free

```
1 counters with mean arrival and service times 1.5 2.0 minutes
Customer00  0.00  0.00  2.02
Customer01  1.62  0.86  3.69
Customer02  4.28  1.56  2.82
Customer03  7.36  0.00  1.02
Customer04  9.42  0.00  0.81
```

And here is an output with two counters.

```
2 counters with mean arrival and service times 1.5 2.0 minutes
Customer00  0.00  0.00  2.54
Customer01  0.87  0.00  2.01
Customer03  2.04  0.84  0.78
Customer04  2.84  0.77  1.93
Customer02  1.78  0.75  2.48
```

You will notice that the customers do not necessarily exit in the same order as they entered. Also, even with two counters, there are times when a queuing delay does occur.

Obviously, you need to experiment with a large number of customers. Besides, printing the results is not particularly useful or meaningful.

Statistical analysis of the simulation

You can collect the data in the process function and then use SciPy's statistical functions. SimPy includes some simple statistical methods, which cannot be explored here.

```
def process(name, time_in, wait_time, time_in_bank):
    wait_array.append(wait_time)
    tib_array.append(time_in_bank)

import scipy as np
wait_array=[]
tib_array=[]
main(1000, 100, [1.5, 2.0, 2])
wt=np.array(wait_array)
tib=np.array(tib_array)
print("%15s" %("Mean", "Variance", "Median", "Max"))
print("%s %6.2f %6.2f %6.2f %6.2f" %("Queue Time",
   wt.mean(), wt.var(), np.median(wait_array), wt.max()))
print("%s %6.2f %6.2f %6.2f %6.2f" %("Time in Bank",
    tib.mean(), tib.var(), np.median(tib_array), tib.max()))
print("The percent of customers who did not have to wait %4.1f"
    %(100 - 100.0*len(wt.nonzero()[0])/len(wt)))
```

You keep appending times to the global Python lists during the simulation. The lists are converted to NumPy arrays and their mean, variance, median and maximum are printed. Of these, the median is not a property of a NumPy array and is computed using a standard NumPy method. With two counters, a customer is handled every minute; hence, it's also interesting to see how many customers did not have to wait. Numpy arrays have a number of useful properties—for example, to get indices of the elements that are not zero.

The results for a run with two counters will look like the following code:

```
2 counters with mean arrival and service times 1.5 2.0 minutes
            Mean  Variance  Median  Max
Queue Time    0.68    1.00    0.00   3.78
Time in Bank  2.56    2.09    2.11   8.48
The percent of customers who did not have to wait: 58.0
```

With a sample of only a hundred, running the simulation again may give you substantially different results. You can experiment with larger sample sizes and explore the many more functions available in SciPy.

With this article, we come to the end of our exploration of SciPy, with the hope that you will find it useful when solving more complex research and educational problems. END

References

- SimPy Tutorial: *http://simpy.sourceforge.net/SimPyDocs/SimPy_Tutorials.html*
- Erlang Distribution: *http://en.wikipedia.org/wiki/Erlang_distribution*
- SciPy Statistics: *http://docs.scipy.org/doc/scipy/reference/stats.html*

By: Anil Seth

The author works as a consultant. You can find him online at http://sethanil.com and reach him via email at seth.anil@gmail.com

S.G. Ganesh

Understanding Increment and Decrement Operators in C

Almost all programmers who work in C based languages use the ++ and - operators. Let's look at some interesting aspects of these operators in C.

Increment/decrement operators are good examples to explain the curt (concise) syntax of C language

```
void strcpy(char *t, const char *s) {
  while(*t++ = *s++);
}
```

In this *strcpy* function the *while* loop continues copying *chars* from source to target, till it reaches the end of the string ('\0' character)

Here is another example; can you guess the output of this code segment?

```
char str[] = "oh", *p = str;
while (*p)
  ++*p++;
printf(str);
```

How is the expression ++*p++ treated? The postfix increment/decrement is part of the operators with highest precedence; the dereference and prefix increment/decrement is at the next precedence level and their associativity is from right to left. So, the expression is treated as: *(++(*(p++)))*. *p* is a pointer to string *"oh"*. In the *while* loop *p++* is executed first, but its effect doesn't take place till the statement execution completes. The result of *p++* (which is still the value of *p*) is dereferenced with *. For the first character it is 'o'. Now prefix ++ increments the character value, so its value is 'p'. The same happens for the next character 'h', which becomes 'i'. Now the end of the character is reached, so the *while* loop terminates. The *printf* now prints the modified string, which is "pi".

We can use ++/-- operators to illustrate the maximum munch rule in C lexical analysis. Simply stated, the lexical analyser (lexer) looks for a match for the longest token. For example, consider the expression *a+b* where *a* and *b* are integer variables. Given this as input in a C program, the lexer might give VARIABLE (a) PLUS (+) VARIABLE(b) as output, which is a valid sequence of tokens for compilation to succeed.

Now how about the expression *a+++b*? It is treated as *a ++ b*, and not *(a + (+b))* and results in a compiler error. How about *a+++b*? It becomes *((a++) + b)*, so it's fine. Expression *a++++b*? It is *((a++) (++b))*, so it is a compiler error. Hmm, how about *a+++++b*? It is *((a++) ++) (+b))* which is again a compiler error. Note that, in all these cases, if we use explicit white-space or parenthesis, we can get valid expressions: *(a+ (+b))*, *((a++) + b)*, *((a++) + (+b))* and *((a++) + (++b))*. If we try adding more +s after this, we cannot get a valid expression for integer arguments.

Today C-based languages (C++, Java, etc) rule the programming world; because of this, almost all programmers are aware of the ++ and -- operators and how powerful they are in expressing operations. What even experienced programmers do not know is that the C family was the first to introduce the ++ and -- operators! Ken Thompson, while implementing the B compiler, noticed that increment operations generate more compact code than adding 1 to a variable and storing it back. In other words, an increment was one instruction while adding 1 and storing it back was more than one instruction.

So he implemented the prefix ++ (and --) operators, and generalised it by adding a postfix version. Interestingly, this short code-generation advantage is still true for recent languages. For example, in Java, integer increment would generate one instruction (*iinc*) compared to the four instructions for adding 1 to the variable and storing it back (*iload_x*, *iconst_1*, *iadd* and *istore_x* instructions)!

About the author

S G Ganesh is a research engineer in Siemens (Corporate Technology). His latest book is "Cracking the C, C++ and Java Interview", published by Tata McGraw-Hill. You can reach him at sgganesh@gmail.com

GO
GREEN IT
LEARN HOW TO
VIRTUALIZE FROM
VERY TALENTED PEOPLE